An Analysis of the EMV Channel Establishment Protocol

C. Brzuska[1], N.P. Smart[2], B. Warinschi[2], and G.J. Watson[2]

[1] School of Computer Science, School of Engineering
Tel Aviv University, Israel.
[2] Dept. Computer Science,
University of Bristol, UK.

Abstract. With over 1.5 billion debit and credit cards in use worldwide, the EMV system (a.k.a. "Chip-and-PIN") has become one of the most important deployed cryptographic protocol suites. Recently, the EMV consortium has decided to upgrade the existing RSA based system with a new system relying on Elliptic Curve Cryptography (ECC). One of the central components of the new system is a protocol that enables a card to establish a secure channel with a card reader. In this paper we provide a security analysis of the proposed protocol, we propose minor changes/clarifications to the "Request for Comments" issued in Nov 2012, and demonstrate that the resulting protocol meets the intended security goals.

The structure of the protocol is one commonly encountered in practice: first run a key-exchange to establish a shared key (which performs authentication and key confirmation), only then use the channel to exchange application messages. Although common in practice, this structure takes the protocol out of the reach of most standard security models for key-exchange. Unfortunately, the only models that can cope with the above structure suffer from some drawbacks that make them unsuitable for our analysis. Our second contribution is to provide new security models for channel establishment protocols. Our models have a more inclusive syntax, are quite general, deal with a realistic notion of authentication (one-sided authentication as required by EMV), and do not suffer from the drawbacks that we identify in prior models.

1 Introduction

The EMV chip-and-pin system is used to secure the majority of the world's credit card and ATM transactions and protects electronic banking in many countries. The current system uses RSA public-key cryptography, combined with DES based symmetric-key cryptography. In the EMV system, bank or credit card customers are issued with a plastic card containing an embedded chip holding various cryptographic keys and which can perform various cryptographic operations. The card is used to communicate with a terminal (typically a point-of-sale terminal in a shop, but other terminals are possible). In addition the card can produce cryptograms for sending on to the banking system for processing. Nonetheless, the cryptographic functionality provided by the card in its first generation incarnation is relatively limited.

As part of a major reworking of the chip-and-pin system, the EMV consortium has decided to replace RSA with ECC based systems and to let the card provide additional cryptographic functionalities. In Nov 2012 EMVCo released a Request-For-Comments [13] about a draft specification for an important sub-protocol within the system; namely a protocol that allows a card to establish a channel with a terminal. Calling the security of these protocols "important" is a serious understatement: the total number of public keys and certificates (1.55 billion as of Q2 2012) deployed in the EMV systems dwarfs the paltry 5.8 million TLS certificates found in [9].

The problem of establishing and implementing secure channels is central to practical uses of cryptography and a superficial look at existing literature would lead one to believe that this is a solved problem. What can be simpler than first running a secure key-exchange protocol and then using the resulting keys to somehow encrypt and authenticate the messages to be sent? Indeed, there are a plethora of works looking at key establishment [3, 4, 7] and a similar number of works looking at how to build secure channels on top of shared keys [1, 5, 15]. However, the traditional key agreement models such as those following the schema set out by Bellare and Rogaway [3] have been shown to be less usefully applicable to deployed protocols. In particular the property that keys derived in secure key-exchange protocols are indistinguishable from truly random keys is often broken in practice by explicit key confirmation steps.

The realization that real-world protocols, like TLS, are therefore outside the reach of the traditional models for key-exchange and channels, has triggered renewed interest in formal models for secure channels [14, 11, 6].

These approaches deal with what is essentially an overlap between the key-exchange part and the secure channel part of a channel-establishment protocol by either modifying the protocol, analyzing the overall protocol monolithically, or developing methods that allow for a modular analysis despite the overlaping phases.

The structure of the EMV protocol for establishing channels follows the recipe described above: during the key-exchange phase itself, the channel is already used before the deployed keys are accepted; and the messages that are sent over the channel are crucial for the security of the overall protocol. Our work can therefore be seen as a continuation of the recent thrust on research on models for channel establishment protocols. Below we describe the state-of-the art for such models, identify some of their weaknesses, and overview our results.

EXISTING MODELS TO SECURE CHANNEL ESTABLISHMENT. There are currently two approaches to studying the combined properties of a key establishment and a secure channel protocol, when there is no clear separation between the two components. Very roughly, the first approach is to relax the security requirement on the keys by demanding that they are sufficiently strong to be used for the primitives that make-up the channel, and then show that the channel security relies only on these primitives. This *modular approach* is explored in [6], where a game-based composition theorem is provided for combining key agreement protocols with other protocols using the previously agreed keys. The approach is shown to work for real world protocols such as TLS.

In this paper, we prefer to avoid the machinery needed to work within this framework and instead concentrate on the approach of Jager et al. [11]. They propose to analyze channel establishment protocols, monolithically, with respect to security models devised for this specific task. The models that they give are tailored for TLS and are not immediately applicable for EMV. Worse, both the original version of the model [11] and a more recent refined version [10] do not seem to appropriately capture the level of security that one would like. In brief, the former model is too strong, to the point that it actually rules as insecure protocols like TLS and the one that we analyze in this paper. The refined version, on the other hand seems to be too weak, as it takes away one of the adversary's abilities, an ability that reflects possible real-world powers.

A bit more in detail, the issue concerns the ability of an adversary to "reveal" a key, and its interaction with how "partnering" is defined. We explain these concepts next. Traditionally, reveal queries model the unintended leakage of session keys from a participant to assure that keys which leak from one session should not affect the security of other sessions. Partnering formalizes the intuition that each session of a protcol should somehow be matched with a single session of some intended partner.

One of the first formulation of partnering relies on "matching" conversations (the outgoing/incoming messages of one sessions are the same as the incoming/outgoing messages of its partner). The requirement is that if a session accepts, then it has had a matching conversation with "the right" partner. Unfortunately, in any protocol where some messages are sent encrypted with the key that is derived, the above requirement cannot hold. An attacker can proceed as follows. When an encrypted message goes on the network, block it, reveal the key that was used for encryption and then send to the recepient a different encryption of the same message, deploying fresh randomness. The two partners will not have a matching conversation, although the protocol will be executed successfully. In Appendix C.2 we describe an attack against entity authentication via matching conversations for TLS when an adversary is permitted to reveal keys as soon as they are derived. We stress that our attack uncovers subtleties in modelling and is not an actual attack on the TLS protocol.

One approach to circumvent this attack is to preclude the adversary from performing such a reveal. This is the approach taken in [10] which only consider that keys can be revealed once the session in which they are derived had accepted. We find this restriction unsatisfactory. If reveals are considered possible, then they should be able to target a key as soon as that key has been derived. In particular, revealing a key that has just been used to perform an encryption should be allowed. A more indepth discussion of weaknesses present in the existent models for channel-establishment protocols is in Appendix C.2.

Our Contribution

In this paper we present a new definitional framework which addresses the problems identified in previous approaches. In particular we present a security model which is particularly tailored to the case of key-exchange

followed by the creation of a secure channel. Our new framework is conceptually simpler than previous models and can be further extended to capture one-sided key agreement followed by composition with a secure channel. Below we highlight some of our contributions and techniques.

MODELS FOR CHANNEL-ESTABLISHMENT PROTOCOLS. For entity authentication, we deal with the realistic case of one-sided authentication. This is demanded by the protocol that we analyze which is inherently one-sided: the card is authenticated, while the terminal is not. We remark that existing models for channel-establishment concentrate on mutual authentication, and the case of one-sided authentication had been considered only sporadically in the key-exchange literature.

We also provide a satisfactory solution to the issue of bad interaction between partnering and reveal queries. Here, we take a different route than [11, 10]. Instead of weakening the adversary, we chose to relax the partnering requirement: we only demand that partners agree on some common session identifier. This approach originates in the work of Bellare, Pointcheval, and Rogaway [2], is quite common in key-exchange literature, and reflects an intuitively appealing level of security.

Finally we model and analyze unlinkability properties of the proposed protocol from EMVCo. One of the design criteria of the protocol is a mild form of unlinkability; an adversary that sees a message flow between a terminal and a card should not be able to link this card's current transaction with a previous transaction from the same card. The protocol aims to ensure this by not transmitting the certificate in the clear, however the proposed protocol also uses a performance optimization in that the card uses a small ephemeral private key. We establish that using a small ephemeral key in this way should be avoided.

PROTOCOL ANALYSIS AND EMV RECOMMENDATIONS. The EMV channel establishment protocol consists of a key exchange phase and an application phase. The key exchange phase is an ECC-based Diffie-Hellman-like protocol with one-sided authentication. We analyze the EMV channel establishment protocol and identify the assumptions under which it can be proved secure with respect to the notion that we put forth. We end this introduction by pointing out a number of recommendations related to the EMV protocol which have been passed to the designers as a result of our analysis:

1. The resulting Diffie–Hellman key should be hashed down to obtain the used symmetric keys. The proposal in [13] says to use a hash function or the x-coordinate of the elliptic curve point as the key derivation function. We consider any choice not using a hash function to be insecure; indeed our security analysis crucially relies on the hash being applied.
2. The resulting keys should be used in a uni-directional manner; thus two keys need to be obtained from the hashing process. This avoids a large number of potential replay attacks on the application layer. If this was not done, the application layer would need to be implemented extremely carefully to thwart these attacks. Having two keys, one for each direction, makes the design of a secure application layer less vulnerable. We have implicitly assumed, as this is not stated in [13], that the resulting secure channel should be secure against adversaries both deleting messages and playing messages out of order; since this is the usual definition of a secure channel.
3. The card ephemeral key a should *not* be selected from the set $\{0,1\}^{32}$. We suggest that it is not restricted in size and instead chosen at random from $\mathbb{F}_q$. If the value a is selected from the set $\{0,1\}^{32}$, then this has a significant effect on security. Not only does it reduce the scheme's ability to achieve unlinkablity. But in addition, when a is selected from a small set an adversary could establish two sessions of one card which share the same key with a single terminal.

2 Scheme

Our presentation follows that in [13], augmented with information obtained from public discussions with the authors of the protocol at several meetings. The basic underlying idea of the protocol is to use a Diffie–Hellman key exchange in which one side (the card) has a static public key. In order to achieve unlinkability the certificate of this public key is not passed in the clear; instead, the card's static Diffie–Hellman key share is randomized

by an additional ephemeral secret. The resulting Diffie–Hellman key is then hashed using a cryptographic hash function; which we will model as a random oracle.

The Diffie–Hellman group used by the protocol is defined over an elliptic curve $G = E(\mathbb{F}_p)$ having group order a prime q. The prime q is a function of an implicit security parameter k, but in practice the group is fixed and so all our results are given in the concrete security setting. Along with the group G a base point $P \in G$ is given.

After the protocol has established secret keys these are used in a secure channel protocol $(\mathsf{SendCh}, \mathsf{ReceiveCh})$. On input an application message m and state st_e, SendCh returns a channel message $\mathfrak{ch}$. On input a channel message $\mathfrak{ch}$ and state st_d, $\mathsf{ReceiveCh}$ returns an application message m. The secure channel protocol is based on a stateful authenticated-encryption (AE) scheme $\mathsf{AE} = \{\mathsf{enc}, \mathsf{dec}\}$. We assume that all plaintext headers used by the secure channel are unauthenticated, implying that no header is sent in clear as part of the AE scheme. The states st_e and st_d here model the fact that in practice sequence numbers are used to ensure that messages are delivered in order, thus the operations are stateful. We assume that the underlying authenticated encryption scheme satisfies the standard properties of indistinguishability under chosen message attack and integrity of ciphertexts for such stateful schemes, assuming the key-agreement scheme has generated a randomly distributed key. See Appendix A for precise definitions of these security notions for a secure channel.

We also assume that there is a public key signature algorithm used to define certificates. In particular each card C has a long term public/private key pair (Q_C, d), where $d \in \mathbb{F}_q$ and $Q_C = dP \in G$. A certificate is a signature/message pair $\mathrm{cert}_C = (\mathsf{sig}_{\mathsf{sk}}(Q_C), Q_C)$ provided by an issuing authority with a public/private key pair $(\mathsf{pk}, \mathsf{sk})$ for some (unspecified) public key signature algorithm $(\mathsf{sig}, \mathsf{ver})$. All that we require of the signature algorithm is that it be existentially unforgeable under a chosen message attack. Again Appendix A gives the precise security definition we will use.

We are now in a position to define the EMV key establishment and secure channel protocol in Figure 1. As well as the components above the protocol makes use of a hash function H which takes elements in the group G and maps them onto a pair of keys for the authenticated encryption scheme. The keys are used to secure the communication in both directions; we propose the use of two keys so that replay attacks are prevented at the level of the protocol as opposed to needing to be dealt with at the application layer.

Note that when we perform our security analysis later we will make use of a session identifier to ensure unique partnering. We shall define the session identifier to equal the pair of keys that are derived.

Card (C)		**Terminal (T)**
$a \xleftarrow{r} \mathbb{F}_q$	$\xrightarrow{A=aQ_C}$	
	$\xleftarrow{E=eP}$	$e \xleftarrow{r} \mathbb{F}_q$
$(\kappa_e^C, \kappa_d^C) = \mathsf{H}(daE)$		$(\kappa_d^T, \kappa_e^T) = \mathsf{H}(eA)$
$(\mathfrak{ch}, st_e^C) = \mathsf{SendCh}_{\kappa_e^C}(\mathrm{cert}_C \Vert Q_C \Vert a; st_e^C)$	$\xrightarrow{\mathfrak{ch}}$	$(\mathrm{cert}_C \Vert Q_C \Vert a, st_d^T) = \mathsf{ReceiveCh}_{\kappa_d^T}(\mathfrak{ch}; st_d^T)$ Check $\mathsf{ver}_{\mathsf{pk}}(\mathrm{cert}_C, Q_C) \stackrel{?}{=} \mathsf{true}$ Check $aQ_C \stackrel{?}{=} A$
$(\mathfrak{ch}_1, st_e^C) = \mathsf{SendCh}_{\kappa_e^C}(m_1; st_e^C)$	$\xrightarrow{\mathfrak{ch}_1}$	$(m_1, st_d^T) = \mathsf{ReceiveCh}_{\kappa_d^T}(\mathfrak{ch}_1; st_d^T)$
$\cdots$		$\cdots$
$(m_2, st_d^C) = \mathsf{ReceiveCh}_{\kappa_d^C}(\mathfrak{ch}_2; st_d^C)$	$\xleftarrow{\mathfrak{ch}_2}$	$(\mathfrak{ch}_2, st_e^T) = \mathsf{SendCh}_{\kappa_e^T}(m_2; st_e^T)$.

Fig. 1. Combined Authenticated Key Agreement Scheme and Secure Channel Protocol. Note that $\kappa_e^C = \kappa_d^T$ and $\kappa_d^C = \kappa_e^T$.

As mention in the introduction the proposal by EMVco [13] suggest that the ephemeral secret[1] a should be small, (less than 2^{32}). They state that this choice is "set to be fit for purpose for blinding a one-off session key". First note that the unlinkability property may be hard to achieve when a is small: Given two public keys Q_1 and Q_2 and the first message of a session aQ_i, there is an obvious sqrt-root attack which determines Q_i when a is small, i.e. an attack which runs in time roughly 2^{16} operations.

More seriously, the security of entity authentication would also be at risk. An adversary can perform an attack which allows two sessions of a card to share the same key with a single terminal. This would break the uniqueness of sessions requirement that will be necessary to achieve security and allowing the possibility of replay attacks to occur. There are other approaches which could prevent this latter issue (in cases where unlinkability is not an issue) but we believe increasing the size of a to be the simplest and offer least chance of implementation errors being introduced. In the rest of the paper we assume that a is chosen from $\mathbb{F}_q$ and therefore our security results apply only to this case.

3 New Security Models

In this section we present our security models for the secure channel establishment and unlinkability. Most of the section is however devoted to the more complex case of modeling secure channel establishment.

PRELIMINARIES. Before giving our new definition we present some preliminary definitions. Let I be the set of participants. Each participant has a distinct ID i, long-term public key pk_i and corresponding secret key sk_i. The protocol description is defined by two efficiently computable stateful algorithms $P = \{\Pi, \mathcal{G}\}$. The algorithm Π defines how honest parties behave and $\mathcal{G}$ is a public/private key pair generation algorithm. Each execution of this algorithm maintains the following state information:

- $st_k \in \{0,1\}^*$ is some state information for the key exchange.
- $\delta \in \{\texttt{derived}, \texttt{accept}, \texttt{reject}, \perp\}$ is current state of the key-exchange (initialised to $\perp$).
- $\rho \in \{\texttt{initiator}, \texttt{responder}\}$ is the role of the participant.
- sid a session identifier.
- pid a partner identifier
- $\kappa = (k^{\rho}_{\mathsf{enc}}, k^{\rho}_{\mathsf{dec}}) \in (\{0,1\}^* \cup \{\perp\})^2$. This is the agreed pair of keys. The order of these keys depends on the role $\rho = \{\texttt{initiator}, \texttt{responder}\}$ and $\kappa = (\perp, \perp)$ unless $\delta = \texttt{derived}$.

A CLASS OF PROTOCOLS. We define a specific class of protocol based on the combination of a key-exchange protocol and a secure channel. The key-exchange may include some steps where messages are sent using the newly established secure channel, i.e. after keys are derived but before they are accepted. We define the honest operation of a participant engaged in such a protocol via a program $\Pi = (\mathsf{KeyExch}, \mathsf{SendCh}, \mathsf{ReceiveCh})$. Here KeyExch, SendCh and ReceiveCh define the respective algorithms for key-exchange, sending a message on the channel and receiving a message from the channel. The syntax of Π is given in Figure 2.

The program Π takes as input a message m and an operation $\texttt{op} \in \{\texttt{SendCh}, \texttt{ReceiveCh}\}$ which the user requests to be performed. An operation request op is only carried out if keys have been accepted ($\delta = \texttt{accept}$), prior to this messages are dealt with appropriately by KeyExch.

The algorithm KeyExch takes as input a message m and a state st_k, and outputs a new message m' followed by six further variables $\texttt{op}_{next}, \texttt{op}_{now}, st_k, \delta, \kappa_e, \kappa_d$, (all initialised as $\perp$). The state st_k is used to manage the internal state of KeyExch. The state δ is set to either $\perp, \texttt{derived}, \texttt{accept}$ or $\texttt{reject}$ and defines the current view of the keys. Once $\delta = \texttt{derived}$, the derived keys $(\kappa_e, \kappa_d) \neq (\perp, \perp)$ are output by KeyExch for use by the secure channel. At this point the program Π initialises the states of the secure channel (st_e, st_d) by calling $\mathsf{initial}(st_k)$. We are now in the key-confirmation phase where key-exchange messages can be sent on the secure channel. The algorithm KeyExch uses the states $\texttt{op}_{now}$ and $\texttt{op}_{next}$ to keep track of when a key-exchange message should be sent on the secure channel and when the next message should be received from the secure

[1] In the EMV draft a is denoted r.

channel, respectively. Finally, once KeyExch outputs $\delta = \texttt{accept}$ the secure channel has been successfully established.

An input $(m, \texttt{op})$ with $\texttt{op} \in \{\texttt{SendCh}, \texttt{ReceiveCh}\}$ is dealt with as follows. The operation SendCh specifies that the input m is an application message (to be sent on the channel) and this should be dealt with by calling SendCh. If a key has not yet been accepted then the program will return $\perp$. The operation ReceiveCh specifies that the input m is a channel message (i.e. a message received from the channel) and this should be dealt with by calling ReceiveCh. If a key has not yet been accepted then the program should forward the message m to KeyExch.

PROGRAM $\Pi(m, \texttt{op})$:
$m' := \perp$
if $\delta = \perp$ or $\texttt{derived}$ and $\texttt{op} \neq \texttt{SendCh}$ **then**
  if $\texttt{op}_{next} = \texttt{ReceiveCh}$ **then** $(m, st_d) \leftarrow \mathsf{ReceiveCh}_{\kappa_d}(m; st_d)$
  $(m', \texttt{op}_{next}, \texttt{op}_{now}, st_k, \delta, \kappa_e, \kappa_d) \leftarrow \mathsf{KeyExch}(m; st_k)$
  if $\delta = \texttt{derived}$ and $\gamma = \mathsf{false}$ **then** $(st_e, st_d) \leftarrow \mathsf{initial}(st_k)$; $\gamma := \mathsf{true}$
  if $\texttt{op}_{now} = \texttt{SendCh}$ **then** $(m', st_e) \leftarrow \mathsf{SendCh}_{\kappa_e}(m'; st_e)$
else if $\delta = \texttt{accept}$ **then**
  if $\texttt{op} = \texttt{SendCh}$ **then** $(m', st_e) \leftarrow \mathsf{SendCh}_{\kappa_e}(m; st_e)$
  if $\texttt{op} = \texttt{ReceiveCh}$ **then** $(m', st_d) \leftarrow \mathsf{ReceiveCh}_{\kappa_d}(m; st_d)$
return m'

Fig. 2. Honest Protocol Execution

EXECUTION MODEL. We consider the standard execution model for key exchange protocols where an adversary $\mathcal{A}$, is assumed to control all communication between participating parties i.e. the adversary can intercept all messages sent and inject any message that he wishes. Let $\Pi^s_{i,j}$ denote the oracle modelling participant $i \in I$ engaged in the protocol described above with participant $j \in I$ in session s. Each oracle $\Pi^s_{i,j}$ runs the program Π and maintains the states of that program instance. The adversary can make the following queries:

- $\mathsf{NewSession}(i, \rho)$: Create a new session for user i with role ρ.
- $\mathsf{Send}(\Pi^s_{i,j}, m, \texttt{op})$: Sends message m to $\Pi^s_{i,j}$ with operation op. As a result $\Pi^s_{i,j}$ will run the program Π on input m and op.
- $\mathsf{Reveal}(\Pi^s_{i,j})$: reveals the current session key κ of $\Pi^s_{i,j}$.
- $\mathsf{Corrupt}(i)$: reveals the long-term private key of i.

PARTNERING AND FRESHNESS. In order to define security for key-exchange protocols it is necessary to define the notion of partnering. Two participants should only establish a shared key if they have been successfully partnered. There are many approaches to defining such a notion. We begin by discussing the concept of *matching conversations*, introduced by Bellare and Rogaway [3] in the context of authenticated key exchange. A participant's conversation can be defined as a transcript of all the messages it receives and sends. As the name suggest, matching conversations defines two participants to be partnered if their transcripts *match*. It is this approach which is followed by Jager et al. [11] in their definition of ACCE. Unfortunately, when protocols use the session key to encrypt messages as part of a key confirmation step, attacks may be possible which violate the requirements of matching conversations[2] (cf. Appendix C.2). Notice however that while the attack described violates the matching conversation property, should perhaps not be considered an attack. The plaintext that was sent by one party reached its intended recepient. We interpret this attack as a limitation of the model: it may rule out as insecure protocols with no obvious weaknesses.

[2] The adversary reveals the key and then uses this to re-encrypt the confirmation message with new randomness. The two transcripts now differ for this message.

Our formulation uses a definition of partnering based on session identifiers [2]. Informally, we declare two oracles partnered if they have already derived keys and i) they both share the same session identifier sid, ii) they derived the same key κ, and iii) one oracle is an initiator and the other a responder. Moreover, to ensure each oracle accepted with only a single partner we also ask that iv) there should exist no other oracle which has derived keys and holds the same session identifier. The intuition is captured by the following defintion.

Definition 1 (Partner). *We say that oracles $\Pi^s_{i,j}$ and $\Pi^t_{j,i}$ holding (κ, sid, pid) and (κ', sid', pid') respectively are partners if they have both derived keys ($\delta = \texttt{derived}$) and the following three conditions hold:*

- *$sid = sid'$, $\kappa = \kappa'$ and $pid = j$ and $pid' = i$.*
- *$\rho^s_i = \texttt{initiator}$ and $\rho^t_j = \texttt{responder}$, or $\rho^s_i = \texttt{responder}$ and $\rho^t_j = \texttt{initiator}$*
- *No oracle besides $\Pi^s_{i,j}$ and $\Pi^t_{j,i}$ that have derived keys, have session identifier sid.*

We make the following remark about a slight difference between our definition and that of [2]. Bellare et al. make a distinction between an oracle accepting and terminating. Accepting defines the event that the session keys have been established but the key confirmation steps are still to follow. An oracle terminates after the key confirmation steps have completed. Once keys are accepted they may be revealed but the key-exchange protocol has yet to terminate. We argue that a key is not "accepted" until after the key confirmation step since this step may fail. As a result, we use the terms derived and accepted, where derived corresponds to Bellare et al.'s accepted and our accepted corresponds to their terminated.

A concept that plays a central role in defining security in two-party protocols is that of "freshness". Intuitively, an oracle is fresh if it has accepted and an adversary had not "tampered" with it in any way, i.e. the adversary has not revealed or corrupted the oracle or its partner. A notion of freshness is necessary when defining security since the security guarantees are only for such oracles. The next definition formalises the concept.

Definition 2 (Fresh). *An oracle $\Pi^s_{i,j}$ is **fresh** if the following three conditions hold:*

1. *$\Pi^s_{i,j}$ has accepted.*
2. *Oracle $\Pi^s_{i,j}$ has not been revealed and user i is not corrupted.*
3. *No partner oracle of $\Pi^s_{i,j}$ has been revealed and no parent of such a oracle has been corrupted.*

3.1 Security Definitions: Two-Sided Authentication Setting

We formulate three levels of security: entity authentication, message authentication and message privacy. The later definitions rely on entity authentication and we start by defining that definition.

ENTITY AUTHENTICATION. We consider that an adversary violates entity authentication if he can get a session to accept even if there is no unique session of its intended partner that has derived the same key. More formally, we wish to verify that there exists no oracle that accepts without a partner oracle. Following on from Definition 1 we again follow the definitions from [2].

First consider the entity authentication experiment $\mathsf{entauth}$ that generates public/private key pairs for each user $i \in I$ (by running $\mathcal{G}$) and returns the public keys to $\mathcal{A}$. The experiment then allows the adversary $\mathcal{A}$ to make the queries $\mathsf{NewSession}(i, \rho)$, $\mathsf{Reveal}(\Pi^s_{i,j})$, $\mathsf{Corrupt}(i)$ as well as $\mathsf{Send}(\Pi^s_{i,j}, m, \mathsf{op})$ with operations $\mathtt{op} \in \{\mathtt{SendCh}, \mathtt{ReceiveCh}\}$. We say that an adversary violates entity authentication (and hence "wins" this experiment) if an oracle accepts but has no uncorrupted partner oracle and define the probability of this to be the adversary's advantage $\mathbf{Adv}^{\mathsf{entauth}}_{\Pi}(\mathcal{A}_{\mathsf{ent}})$.

Definition 3 (Entity Authentication (EA)). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a (t, ϵ_{EA})-**secure EA protocol** if for all adversaries $\mathcal{A}_{\mathsf{ent}}$ running in time at most t, $\mathbf{Adv}^{\mathsf{entauth}}_{\Pi}(\mathcal{A}_{\mathsf{ent}}) \leq \epsilon_{EA}$.*

To define the security experiments for message authentication and privacy we shall make use of the following notation for lists maintained for each $\Pi^s_{i,j}$ as follows:

– Application messages sent $L^{app|sen}_{i,j,s}$, i.e. the list of all messages m input to $\mathsf{Send}(\Pi^s_{i,j}, m, \mathtt{SendCh})$.
– Channel messages sent $L^{ch|sen}_{i,j,s}$, i.e. the list of all outputs from $\mathsf{Send}(\Pi^s_{i,j}, m, \mathtt{SendCh})$.
– Channel messages received $L^{ch|rec}_{i,j,s}$, i.e. the list of all messages m input to $\mathsf{Send}(\Pi^s_{i,j}, m, \mathtt{ReceiveCh})$.
– Application messages received $L^{app|rec}_{i,j,s}$, i.e. the list of all outputs from $\mathsf{Send}(\Pi^s_{i,j}, m, \mathtt{ReceiveCh})$.

MESSAGE AUTHENTICATION. We now turn our attention to message authentication. Here we wish to ensure the integrity and authenticity of all messages sent over the channel. For any two partner oracles $\Pi^s_{i,j}$ and $\Pi^t_{j,i}$, the oracle $\Pi^s_{i,j}$ should only *successfully* receive messages which were output by $\Pi^t_{j,i}$ and vice versa. In the definition which follows we formalise the intuition above by requiring that for any oracle $\Pi^s_{i,j}$ with partner $\Pi^t_{j,i}$, the following holds $\mathsf{Prefix}(L^{app|rec}_{i,j,s}, L^{app|sen}_{j,i,t}) = \mathsf{true}$, where $\mathsf{Prefix}(X, Y)$ is the function which outputs true if X is a prefix of Y (provided not empty) and false otherwise.

Consider the authentication experiment auth that generates public/private key pairs for each user $i \in I$ (by running $\mathcal{G}$) and returns the public keys to $\mathcal{A}$. The adversary is permitted to make the queries $\mathsf{NewSession}(i, \rho)$, $\mathsf{Reveal}(\Pi^s_{i,j})$, $\mathsf{Corrupt}(i)$ as well as $\mathsf{Send}(\Pi^s_{i,j}, m, \mathsf{op})$ with operations $\mathsf{op} \in \{\mathtt{SendCh}, \mathtt{ReceiveCh}\}$. On the $\mathsf{Send}(\Pi^s_{i,j}, m, \mathsf{op})$ query, the game behaves as in Figure 3(a).

For the session matching, we consider the notion of partnering as specified in Definition 1. The notion of freshness that we use in the following definition is according to Definition 2.

We define the following game $\mathbf{Exec}^{\mathsf{auth}}_{\Pi}(\mathcal{A})$ between an adversary $\mathcal{A}$ and challenger $\mathcal{C}$:

1. The challenger $\mathcal{C}$ generates public/private key pairs for each user $i \in I$ (by running $\mathcal{G}$) and returns the public keys to $\mathcal{A}$.[3]
2. Adversary $\mathcal{A}$ is allowed to make as many NewSession, Reveal, Corrupt, Send queries as it likes.
3. The adversary stops with no output.

We say that an adversary $\mathcal{A}$ wins the game if there exists a fresh oracle $\Pi^s_{i,j}$ with partner $\Pi^t_{j,i}$ such that the list $L^{app|rec}_{i,j,s}$ is not a prefix of $L^{app|sen}_{j,i,t}$.

We define the adversary's advantage as:

$$\mathbf{Adv}^{\mathsf{auth}}_{\Pi}(\mathcal{A}) = \Pr[\mathsf{Prefix}(L^{app|rec}_{i,j,s}, L^{app|sen}_{j,i,t}) = \mathsf{false} : \text{for some fresh } \Pi^s_{i,j}].$$

Definition 4 (Message Authenticity (MA)). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a (t, ϵ_{MA})-**secure MA protocol** if for all adversaries $\mathcal{A}_{\mathsf{auth}}$ running in time at most t, $\mathbf{Adv}^{\mathsf{auth}}_{\Pi}(\mathcal{A}_{\mathsf{auth}}) \leq \epsilon_{MA}$.*

MESSAGE PRIVACY. Next we consider the notion of message privacy. Our definition follows the standard indistinguishability paradigm. The adversary should not be able to determine which set of message $\{m_{01}, m_{02}, m_{03}, ...\}$ and $\{m_{11}, m_{12}, m_{13}, ...\}$ has been transmitted on the secure channel.

The message privacy experiment priv initializes the states as in the authentication experiment auth, except that each session now also holds a random secret bit $b^s_{i,j}$. As before, the adversary can make the queries $\mathsf{NewSession}(i, \rho)$, $\mathsf{Reveal}(\Pi^s_{i,j})$, $\mathsf{Corrupt}(i)$. In addition, we introduce a left-right version of $\mathsf{Send}(\Pi^s_{i,j}, m, \mathsf{op})$ which we use to model message privacy. Specifically, query $\mathsf{SendLR}(\Pi^s_{i,j}, m_0, m_1, \mathsf{op})$ takes as input two messages (m_0, m_1) and returns $\mathsf{Send}(\Pi^s_{i,j}, m_{b^s_{i,j}}, \mathsf{op})$. When $\mathsf{op} \neq \mathtt{SendCh}$ we require that these two message are equal, $(\mathsf{SendLR}(\Pi^s_{i,j}, m, m, \mathsf{op}) = \mathsf{Send}(\Pi^s_{i,j}, m, \mathsf{op}))$. As before, two sessions are considered partners by Definition 1. On the $\mathsf{SendLR}(\Pi^s_{i,j}, m_0, m_1, \mathsf{op})$ query, the game behaves as in Figure 3(b).

We define the following game $\mathbf{Exec}^{\mathsf{priv}}_{\Pi}(\mathcal{A})$ between an adversary $\mathcal{A}$ and challenger $\mathcal{C}$:

1. The challenger $\mathcal{C}$, generates public/private key pairs for each user $i \in I$ (by running $\mathcal{G}$) and returns the public keys to $\mathcal{A}$.[4]

[3] Note that in the scheme considered in this paper, public keys of cards are not actually made public to $\mathcal{A}$ but are sent in encrypted form during the confirmation step.

[4] Note that in the scheme considered in this paper, public keys of cards are not actually made public to $\mathcal{A}$ but are sent in encrypted form during the confirmation step.

$\mathsf{Send}(\Pi^s_{i,j}, m, \mathsf{op})$:
$m' \leftarrow \Pi^s_{i,j}(m, \mathsf{op})$
if $\delta = \mathtt{accept}$ and $\mathsf{op} = \mathtt{SendCh}$ **then**
  $L^{app|sen}_{i,j,s} \leftarrow L^{app|sen}_{i,j,s} \| m$
  $L^{ch|sen}_{i,j,s} \leftarrow L^{ch|sen}_{i,j,s} \| m'$
else if $\delta = \mathtt{accept}$ and $\mathsf{op} = \mathtt{ReceiveCh}$ **then**
  if $m' \neq \perp$ **then** $L^{app|rec}_{i,j,s} \leftarrow L^{app|rec}_{i,j,s} \| m'$
return m'

(a) Send query for auth game.

$\mathsf{SendLR}(\Pi^s_{i,j}, m_0, m_1, \mathsf{op})$
if $\delta = \mathtt{accept}$ and $\mathsf{op} = \mathtt{SendCh}$ **then**
  $m' \leftarrow \Pi^s_{i,j}(m_{b^s_{i,j}}, \mathtt{SendCh})$
  $L^{app|sen}_{i,j,s} \leftarrow L^{app|sen}_{i,j,s} \| m_{b^s_{i,j}}$
  $L^{ch|sen}_{i,j,s} \leftarrow L^{ch|sen}_{i,j,s} \| m'$
else if $m_0 \neq m_1$ **then** $m' := \perp$
else
  $m' \leftarrow \Pi^s_{i,j}(m_0, \mathsf{op})$
  if $\delta = \mathtt{accept}$ and $\mathsf{op} = \mathtt{ReceiveCh}$ **then**
    if $m' \neq \perp$ and $\Pi^s_{i,j}$ has a partner $\Pi^t_{j,i}$ **then**
      $L^{app|rec}_{i,j,s} \leftarrow L^{app|rec}_{i,j,s} \| m'$
      $L^{ch|rec}_{i,j,s} \leftarrow L^{ch|rec}_{i,j,s} \| m_0$
      if $\mathsf{Prefix}(L^{ch|rec}_{i,j,s}, L^{ch|sen}_{j,i,t}) = \mathsf{true}$ **then** $m' := \emptyset$
return m'

(b) SendLR query for priv game.

Fig. 3. The Send (resp. SendLR) query for the auth (resp. priv) games

2. Adversary $\mathcal{A}$ is allowed to make as many NewSession, Reveal, Corrupt, SendLR queries as it likes.
3. Finally $\mathcal{A}$ outputs a tuple (i, j, s, b').

We say the adversary $\mathcal{A}$ wins if its output $b' = b^s_{i,j}$ and $\Pi^s_{i,j}$ is fresh. In this case the output of $\mathbf{Exec}^{\mathsf{priv}}_{\Pi}(\mathcal{A})$ is set to 1. Otherwise the output is 0. Formally we define the advantage of $\mathcal{A}$ as

$$\mathbf{Adv}^{\mathsf{priv}}_{\Pi}(\mathcal{A}) = |\Pr[\mathbf{Exec}^{\mathsf{priv}}_{\Pi}(\mathcal{A}) = 1] - 1/2| = |\Pr[b' = b^s_{i,j}] - 1/2|.$$

Definition 5 (Message Privacy (MP)). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a (t, ϵ_{MP})-**secure MP protocol** if for all adversaries $\mathcal{A}_{\mathsf{priv}}$ running in time at most t, $\mathbf{Adv}^{\mathsf{priv}}_{\Pi}(\mathcal{A}_{\mathsf{priv}}) \leq \epsilon_{MP}$.*

We call a channel establishment protocol secure if it satisfies all of the three notions above. We call the resulting notion EAMAP for obvious reasons.

Definition 6 (EAMAP). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a (t, ϵ)-**secure EAMAP protocol** if it is a (t, ϵ)-secure EA protocol, a (t, ϵ)-secure MA protocol and a (t, ϵ)-secure MP protocol.*

Remark 1. Our definitions are with respect to the specific type of protocol construction defined in Figure 2. We note however, that our notions can be extended to more general classes of protocols by simply placing fewer restrictions on the Send queries.

Remark 2. Our mechanism of defining message authentication by requiring that the list of messages received by a party is a prefix of the list of the messages sent by its partner is quite flexible. By appropriately modifying this requirement one can also capture more relaxed notions e.g. where packet dropping or reordering is allowed. Furthermore, we expect that with appropriate restrictions this mechanism can also be adapted to deal with fragmentation. This is a common feature of many secure/authenticated channels in practice and has been formally studied by Boldyreva et al. [5], but is not relevant for EMV.

3.2 Security Definitions: One-Sided Authentication Setting

The above security definitions enforce mutual authentication, yet in many scenarios of practical concern only one party needs to be authenticated. For example, the protocol we consider requires authentication of the credit card but does not authenticate the communicating terminal. To model this situation we split our set of participants I in two. Let C be the set of authenticated participants (the cards) and let T by the set of unauthenticated participants (the terminals), where unauthenticated participants do not hold a long-term private/public key pair.

This formalisation is the same as that of registered and unregistered users in [14]. We say authenticated participants are always initiators and unauthenticated are always responders. As a result of this change we must alter our previous security definitions for entity authentication, message authentication, message privacy and their combination (EAMAP) to consider a one-sided protocol.

ONE-SIDED ENTITY AUTHENTICATION. In the one-sided setting a terminal $j \in T$ wishes to authenticate a card $i \in C$ and establish a key (additionally a secure channel) with this card. Since all $j \in T$ have no long-term secret then it would always be possible for an adversary to impersonate an unauthenticated participant and establish a session with a real card. We need only aim to ensure that a genuine card session is authenticated to an unauthenticated terminal.

Recall the definition of partnering (cf. Definition 1). We define the notion of *one-sided* partnering. The definition that we provide is a stronger version than the natural counterpart of Definition 1. First, we informally describe the notion, formalize it, and then discuss the subtlety involved. A card oracle and a terminal oracle are now said to be os-partners is they both accepted and share the same session identifier sid and key κ, and the card oracle is an initiator and the terminal oracle is a responder. Moreover, to ensure that each genuine terminal oracle that accepts has a single partner we require that for every terminal oracle that accepts, there exists a card oracle which has accepted.

Definition 7 (OS-Partner). *For $i \in C$ **and** $j \in T$, we say that oracles $\Pi^s_{i,j}$ and $\Pi^t_{j,i}$ are **os-partners** if both accept holding (κ, sid, pid) and (κ', sid', pid') respectively and the following three conditions hold:*

- *$sid = sid'$, $\kappa = \kappa'$ and $pid = j$ and $pid' = i$.*
- *$\rho^s_i = \texttt{initiator}$ and $\rho^t_j = \texttt{responder}$.*
- ***if** $\Pi^t_{j,i}$ accepts with session identifier sid **then** there exists a unique $\Pi^s_{i,j}$ which accepts with session identifier sid.*

The main difference between the definition above and the natural restriction of Definition 1 to one-sided partnering is that we consider the partnering guarantees at the moment when oracles accept rather than when keys are derived. In particular this strengthening guarantees that oracles are only partnered *after* they have confirmed the key and accepted to use it to send channel messages. Notice that a similar strengthening does not work for the two-sided case since in this situation one oracle always accepts before the other. An adversary could always ensure that the party that sends the last message of the protocol terminates (and accepts) whereas there would be no corresponding accepting partner. In the one-sided case a terminal will always accept after a card oracle has accepted.

We consider an adversary that violates one-sided entity authentication if he can get a terminal session to accept if there is no unique session of its intended os-partner that has derived the same key. More formally, define the os-entauth experiment in a similar fashion to before but now say that an adversary violates one-sided entity authentication (and hence "wins" this experiment) if an oracle $\Pi^t_{j,i}$ with $j \in T$ accepts but has no uncorrupted *os-partner* oracle. The probability of this event is again defined to be the adversary's advantage $\mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{A}_{\mathsf{ent}})$.

Definition 8 (One-Sided Entity Authentication (OS-EA)). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a (t, ϵ_{EA})-**secure OS-EA protocol** if for all adversaries $\mathcal{A}_{\mathsf{ent}}$ running in time at most t, $\mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{A}_{\mathsf{ent}}) \leq \epsilon_{EA}$.*

ONE-SIDED MESSAGE AUTHENTICITY AND PRIVACY. In order to adapt the definitions of message authenticity and privacy we must consider a one-sided version of freshness. The reason behind this again being that we wish to discount the trivial attack when the adversary impersonates an unauthenticated terminal $j \in T$. A card oracle is defined to be OS-fresh if it has accepted, has not been revealed or corrupted and it is partnered with a genuine terminal oracle. A terminal oracle is defined to be OS-fresh if it has accepted, has not been revealed and it is partnered with a card oracle that has not been revealed or corrupted. We formalise one-sided freshness as follows:

Definition 9 (One-Sided Fresh). *An oracle $\Pi^s_{i,j}$ where $i \in I$ and $j \in I$, is **OS-fresh** if the following six conditions hold:*

1. *Either $i \in C$ and $j \in T$, or $i \in T$ and $j \in C$, i.e. at least one is an authenticated participant.*
2. *$\Pi^s_{i,j}$ has accepted.*
3. *Oracle $\Pi^s_{i,j}$ has not been revealed.*
4. *If $i \in C$ then it is uncorrupted.*
5. *If $i \in C$ then $\Pi^s_{i,j}$ has a partner $\Pi^t_{j,i}$.*
6. *No partner oracle of $\Pi^s_{i,j}$ has been revealed and no parent of such a oracle has been corrupted if they are an authenticated participant.*

Using the above we can alter our previous experiments of auth and priv by requiring that the winning oracle is OS-fresh. We therefore obtain one-sided versions os-auth and os-priv, respectively.

Definition 10 (OS-MA/OS-MP). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a (t, ϵ)-**secure OS-MA protocol** (or OS-MP resp.) if for all adversaries $\mathcal{A}$ running in time at most t, $\mathbf{Adv}^{\mathsf{os\text{-}auth}}_{\Pi}(\mathcal{A}) \leq \epsilon$ (or $\mathbf{Adv}^{\mathsf{os\text{-}priv}}_{\Pi}(\mathcal{A}) \leq \epsilon$ resp.).*

We call a channel establishment protocol with one-sided authentication secure if it satisfies all three of the notions above.

Definition 11 (OS-EAMAP). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a (t, ϵ)-**secure OS-EAMAP protocol** if it is a (t, ϵ)-secure OS-EA protocol, a (t, ϵ)-secure OS-MA protocol and a (t, ϵ)-secure OS-MP protocol.*

3.3 Security Definitions: Unlinkability

A further property that the EMVCo protocol aims to achieve is unlinkability. This means that it should be hard for an adversary to determine when two particular sessions involve the same card. Goldberg et al. [8] define a related notion of anonymity and unlinkability. They aim to prove a scheme secure if an authenticated party remains anonymous to its unauthenticated partner and hence call this *internal anonymity*. Here we are concerned with eavesdroppers external to the execution and hence define a new notion for *external unlinkability*.

We define this security property in terms of the game $\mathbf{Exec}^{\mathsf{unlink}}_{\Pi}(\mathcal{A})$ between adversary $\mathcal{A}$ and challenger $\mathcal{C}$. Informally, the adversary is able to interact with the card and terminal much as in the key agreement game. At some point the adversary halts the first part of his game, and outputs two card identities on which it wishes to be challenged. The challenger then picks one of these two identities and passes to the advsarsary new oracles (i.e. card/teminal session) with respect to the chosen identity. The adversary can then make additional queries, bar Reveal or Corrupt queries on the two test oracles. At the end of the experiment the adversary needs to output which identity the challenger selected. More formally the game is defined as follows:

1. The challenger $\mathcal{C}$, generates public/private key pairs for each user $i \in C$ (by running $\mathcal{G}$) and returns the public keys to $\mathcal{A}$.
2. Adversary $\mathcal{A}$ is allowed to make as many NewSession, Reveal, Corrupt, Send queries as it likes.
3. At some point $\mathcal{A}$ outputs two identities $i_0 \in C$ and $i_1 \in C$.
4. The challenger then chooses a bit $b \xleftarrow{r} \{0, 1\}$ and creates new oracles $\mathcal{O}_C = \Pi^s_{i_b,j}$ and $\mathcal{O}_T = \Pi^t_{j,i_b}$ (for some $j \in T$), by calling NewSession$(i_b, \texttt{initiator})$ and NewSession$(j, \texttt{responder})$.
5. Adversary $\mathcal{A}$ then continues making queries NewSession, Reveal, Corrupt, Send. However, $\mathcal{A}$ is allowed to query oracles $\mathcal{O}_C$ and $\mathcal{O}_T$ only with the Send query.
6. Eventually $\mathcal{A}$ stops and outputs a bit b'.

We say the adversary $\mathcal{A}$ wins if its output $b' = b$ and $\mathcal{O}_C$ and $\mathcal{O}_T$ are OS-partners. In this case the output of $\mathbf{Exec}^{\mathsf{unlink}}_{\Pi}(\mathcal{A})$ is set to one, otherwise the output is zero. Formally we define the advantage of $\mathcal{A}$ as

$$\mathbf{Adv}^{\mathsf{unlink}}_{\Pi}(\mathcal{A}) = |\Pr[\mathbf{Exec}^{\mathsf{unlink}}_{\Pi}(\mathcal{A}) = 1] - 1/2| = |\Pr[b' = b] - 1/2|.$$

Definition 12 (Unlinkability). *A protocol $(\Pi, \mathcal{G})$ is $(t, \epsilon_{\mathsf{unlink}})$-unlinkable, if for all adversaries $\mathcal{A}$ running in time t, $\mathbf{Adv}^{\mathsf{unlink}}_{\Pi}(\mathcal{A}) \leq \epsilon_{\mathsf{unlink}}$.*

4 Main Security Theorems

In this section we state our main security results, and in particular clarify the assumptions under which the EMV channel-establishment protocol is secure. Security of the protocol depends on the signature scheme that is used to produce the certificates and on various assumptions on the group that underlies the scheme. We provide formal definitions of the assumptions in the Appendix.

Theorem 1. *If the Gap-DH problem is hard, the CDH problem is hard,* $\mathsf{AE} = (\mathsf{enc}, \mathsf{dec})$ *is an ind-sfcca secure and int-sfctxt secure authenticated encryption scheme, and the signature scheme* $(\mathsf{sig}, \mathsf{ver})$ *used to produce card certificates is EUF-CMA, then the EMV protocol* Π *in Figure 1 is secure in the sense of OS-EAMAP. In particular we have*

- *If there exists an adversary* $\mathcal{A}$ *running in time at most* t *against the entity authentication property of OS-EAMAP security then there are adversaries* $\mathcal{B}$,$\mathcal{C}, \mathcal{D}, \mathcal{E}$, *such that*

$$\begin{aligned}\mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{A}) \leq \mathbf{Adv}^{\mathsf{eufcma}}_{(\mathsf{sig},\mathsf{ver})}(\mathcal{B}) + n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}^{\mathsf{Gap\text{-}DH}}_{E(\mathbb{F}_p)}(\mathcal{C}) \\ + n_S \cdot n_C \cdot \mathbf{Adv}^{\mathsf{intsfctxt}-0}_{\mathsf{AE}}(\mathcal{D}) + n_C^2 \cdot \mathbf{Adv}^{\mathsf{CDH}}_{E(\mathbb{F}_p)}(\mathcal{E}),\end{aligned}$$

  where $\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}$ *each run in time* $t + O(\mu)$ *where* μ *is total number of bits queried.*
- *If there exists an adversary* $\mathcal{A}$ *running in time at most* t *against the message authentication property of OS-EAMAP security then there are adversaries* $\mathcal{B}$, $\mathcal{C}$ *and* $\mathcal{D}$, *such that*

$$\mathbf{Adv}^{\mathsf{os\text{-}auth}}_{\Pi}(\mathcal{A}) \leq n_S \cdot (n_C + n_T) \cdot \mathbf{Adv}^{\mathsf{intsfctxt}}_{\mathsf{AE}}(\mathcal{D}) + n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}^{\mathsf{Gap\text{-}DH}}_{E(\mathbb{F}_p)}(\mathcal{C}) + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}),$$

  where $\mathcal{B}$ *runs in time* t *and,* $\mathcal{C}$ *and* $\mathcal{D}$ *each run in time* $t + O(\mu)$ *where* μ *is total number of bits queried.*
- *If there exists an adversary* $\mathcal{A}$ *against the message privacy property of OS-EAMAP security then there are adversaries* $\mathcal{B}$, $\mathcal{C}$ *and* $\mathcal{D}$, *such that*

$$\mathbf{Adv}^{\mathsf{os\text{-}priv}}_{\Pi}(\mathcal{A}) \leq n_S \cdot (n_C + n_T) \cdot \mathbf{Adv}^{\mathsf{indsfcca}}_{\mathsf{AE}}(\mathcal{D}) + n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}^{\mathsf{Gap\text{-}DH}}_{E(\mathbb{F}_p),}(\mathcal{C}) + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}),$$

  where $\mathcal{B}$ *runs in time* t *and,* $\mathcal{C}$ *and* $\mathcal{D}$ *each run in time* $t + O(\mu)$ *where* μ *is total number of bits queried.*

where n_C *is the number of cards in the system,* n_T *the number of terminals,* n_S *the number of sessions and* $|h|$ *is the output size of the hash function.*

The proof of this theorem is given in Appendix D. Note that $\mathsf{intsfctxt}{-}0$ defines security for an adversary against $\mathsf{intsfctxt}$ that is permitted no encryption oracle queries.

Finally, we present our theorem for the unlinkability of the protocol:

Theorem 2. *If the gap-DH problem is hard and* $\mathsf{AE} = (\mathsf{enc}, \mathsf{dec})$ *is an ind-sfcca secure authenticated-encryption scheme, then* Π *is secure in the sense of* unlink*; in particular we have*

$$\mathbf{Adv}^{\mathsf{unlink}}_{\Pi}(\mathcal{A}) \leq n_C^2 \cdot \left(\mathbf{Adv}^{\mathsf{indsfcca}}_{\mathsf{AE}}(\mathcal{C}) + n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}^{\mathsf{Gap\text{-}DH}}_{E(\mathbb{F}_p)}(\mathcal{B}) \right)$$

where, again, n_C *is the number of cards in the system and* $|h|$ *is the output size of the hash function.*

The proof of this theorem is given in Appendix E. Note that if a were instead chosen to be of size 2^{32} (as suggested by the RFC) our security analysis would show only 16-bits of security. We refer the reader to the proof for further details.

5 Acknowledgements

This work was support in part by ERC Advanced Grant ERC-2010-AdG-267188-CRIPTO. The second author was also partially supported by a Royal Society Wolfson Merit Award. Research supported in part by the Israel Ministry of Science and Technology (grant 3-9094) and by the Israel Science Foundation (grant 1155/11 and grant 1076/11).

References

1. Mihir Bellare, Tadayoshi Kohno, and Chanathip Namprempre. Breaking and provably repairing the SSH authenticated encryption scheme: A case study of the encode-then-encrypt-and-mac paradigm. *ACM Trans. Inf. Syst. Secur.*, 7(2):206–241, 2004.
2. Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In Bart Preneel, editor, *EUROCRYPT*, volume 1807 of *Lecture Notes in Computer Science*, pages 139–155. Springer, 2000.
3. Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *CRYPTO*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249. Springer, 1993.
4. Simon Blake-Wilson, Don Johnson, and Alfred Menezes. Key agreement protocols and their security analysis. In Michael Darnell, editor, *IMA Int. Conf.*, volume 1355 of *Lecture Notes in Computer Science*, pages 30–45. Springer, 1997.
5. Alexandra Boldyreva, Jean Paul Degabriele, Kenneth G. Paterson, and Martijn Stam. Security of symmetric encryption in the presence of ciphertext fragmentation. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 682–699. Springer, 2012.
6. Christina Brzuska, Marc Fischlin, Nigel P. Smart, Bogdan Warinschi, and Stephen C. Williams. Less is more: Relaxed yet composable security notions for key exchange. *IACR Cryptology ePrint Archive*, 2012:242, 2012.
7. Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *EUROCRYPT*, volume 2045 of *Lecture Notes in Computer Science*, pages 453–474. Springer, 2001.
8. Ian Goldberg, Douglas Stebila, and Berkant Ustaoglu. Anonymity and one-way authentication in key exchange protocols. *Designs, Codes and Cryptography*, 2012. Online first; print version to appear.
9. Nadia Heninger, Zakir Durumeric, Eric Wustrow, and J.Alex Halderman. Mining your Ps and Qs: Detection of widespread weak keys in network devices. In *USENIX Security Symposium – 2012*, pages 205–220, 2012.
10. Tibor Jager, Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of TLS-DHE in the standard model. *IACR Cryptology ePrint Archive*, 2011:219, 2011.
11. Tibor Jager, Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of TLS-DHE in the standard model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO*, volume 7417 of *Lecture Notes in Computer Science*, pages 273–293. Springer, 2012.
12. Caroline Kudla and Kenneth G. Paterson. Modular security proofs for key agreement protocols. In Bimal K. Roy, editor, *ASIACRYPT*, volume 3788 of *Lecture Notes in Computer Science*, pages 549–565. Springer, 2005.
13. EMVCo LLC. EMV ECC key establishment protocols. *http://www.emvco.com/specifications.aspx?id=243*, 2012.
14. Paul Morrissey, Nigel P. Smart, and Bogdan Warinschi. The TLS handshake protocol: A modular analysis. *J. Cryptology*, 23(2):187–223, 2010.
15. Kenneth G. Paterson, Thomas Ristenpart, and Thomas Shrimpton. Tag size does matter: Attacks and proofs for the tls record protocol. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT*, volume 7073 of *Lecture Notes in Computer Science*, pages 372–389. Springer, 2011.
16. John Pollard. Monte Carlo methods for index computation mod p. *Mathematics of Computation*, 32:918–924, 1978.

A Basic Security Definitions

The underlying authenticated encryption scheme we assume satisfies the following two properties which are variants of the stateful security models of Bellare et al. [1] and Paterson et al. [15].

An adversary against a stateful encryption scheme needs to be given the capability to progress the scheme's state without trivially winning the security experiment. It is for this reason that there is a subtle difference between the standard notions of IND-CCA and INT-CTXT, and their stateful versions IND-sfCCA and INT-sfCTXT. An adversary against IND-sfCCA and INT-sfCTXT security is permitted to query the decryption oracle with an output from the encryption oracle in order to progress the state but the output of this query should *not* be returned to the adversary (to avoid the trivial attack).

In order to match with our security definitions of auth and priv we alter the previous definitions of IND-sfCCA and INT-sfCTXT [1] to now compare the lists L^e and L^d. Here L^e is the list of all ciphertexts output

by the encryption oracle and L^d is the list of all ciphertexts *successfully* decrypted by the decryption oracle. In order to prevent an adversary trivially winning he is not permitted to see the output of the decryption oracle if L^d is a prefix of L^e, i.e. if all ciphertexts decrypted so far were output by the encryption oracle.

Definition 13 (IND-sfCCA). *Consider the authenticated-encryption scheme* $\mathsf{AE} = \{\mathsf{enc}_\kappa, \mathsf{dec}_\kappa\}$. *Let* $\mathcal{A}$ *be an adversary with access to a left-or-right encryption oracle* $\mathsf{enc}_\kappa(h, \mathsf{LR}_b(m_0, m_1); st_e)$ *and a decryption oracle* $\mathsf{dec}_\kappa(h, c; st_d)$. *It is mandated that any two messages queried to* $\mathsf{enc}_\kappa(h, \mathsf{LR}_b(m_0, m_1); st_e)$ *have equal length. We define an experiment as follows:*

$\mathbf{Exec}^{\mathsf{indsfcca}-b}_{\mathsf{AE}}(\mathcal{A})$
$\kappa \xleftarrow{r} \{0,1\}^k$, $L := \emptyset$,
$st_e := \emptyset$ *and* $st_d := \emptyset$
Run $\mathcal{A}^{\mathsf{enc}_\kappa, \mathsf{dec}_\kappa}$
  Reply to $\mathsf{enc}_\kappa(h, \mathsf{LR}_b(m_0, m_1); st_e)$ *as follows:*
    $(c, st_e) \xleftarrow{r} \mathsf{enc}_\kappa(h, m_b; st_e)$
    $L^e \leftarrow L^e \cup c$; $\mathcal{A} \Leftarrow c$
  Reply to $\mathsf{dec}_\kappa(h, c; st_d)$ *as follows:*
    $(m, st_d) \xleftarrow{r} \mathsf{dec}_\kappa(h, c; st_d)$
    if $m \neq \perp$ ***then***
      $L^d \leftarrow L^d \cup c$
      if $\mathsf{Prefix}(L^d, L^e) = \mathsf{false}$ ***then*** $\mathcal{A} \Leftarrow m$
Until $\mathcal{A}$ *returns a bit* b'
return b'

The attacker wins when $b' = b$, *and his advantage is defined as*

$$\mathbf{Adv}^{\mathsf{indsfcca}}_{\mathsf{AE}}(\mathcal{A}) = \Pr[\mathbf{Exec}^{\mathsf{indsfcca}-1}_{\mathsf{AE}}(\mathcal{A}) = 1] - \Pr[\mathbf{Exec}^{\mathsf{indsfcca}-0}_{\mathsf{AE}}(\mathcal{A}) = 1].$$

INT-sfCTXT is defined in a similar way. In addition we define the related notion $\mathsf{intsfctxt}{-}0$ which considers an adversary $\mathcal{A}$ against $\mathsf{intsfctxt}$ that is permitted no encryption oracle queries.

Definition 14 (INT-sfCTXT). *Consider the authenticated-encryption scheme* $\mathsf{AE} = \{\mathsf{enc}_\kappa, \mathsf{dec}_\kappa\}$. *Let* $\mathcal{A}$ *be an adversary that has access to the oracles* $\mathsf{enc}_\kappa(h, m; st_e)$ *and* $\mathsf{dec}_\kappa(h, c; st_d)$. *We define an experiment as follows:*

$\mathbf{Exec}^{\mathsf{intsfctxt}}_{\mathsf{AE}}(\mathcal{A})$
$\kappa \xleftarrow{r} \{0,1\}^k$, $L := \emptyset$, $d := 0$,
$st_e := \emptyset$ *and* $st_d := \emptyset$
Run $\mathcal{A}^{\mathsf{enc}_\kappa, \mathsf{dec}_\kappa}$
  Reply to $\mathsf{enc}_\kappa(h, m; st_e)$ *as follows:*
    $(c, st_e) \xleftarrow{r} \mathsf{enc}_\kappa(h, m; st_e)$
    $L^e \leftarrow L^e \cup c$; $\mathcal{A} \Leftarrow c$
  Reply to $\mathsf{dec}_\kappa(h, c; st_d)$ *as follows:*
    $(m, st_d) \xleftarrow{r} \mathsf{dec}_\kappa(h, c; st_d)$
    if $m \neq \perp$ ***then***
      $L^d \leftarrow L^d \cup c$; $\mathcal{A} \Leftarrow 1$
      if $\mathsf{Prefix}(L^d, L^e) = \mathsf{false}$ ***then*** $d := 1$
    else $\mathcal{A} \Leftarrow 0$
Until $\mathcal{A}$ *halts*
return d

The advantage $\mathbf{Adv}^{\mathsf{intsfctxt}}_{\mathsf{AE}}(\mathcal{A})$ *of an adversary is defined as*

$$\mathbf{Adv}^{\mathsf{intsfctxt}}_{\mathsf{AE}}(\mathcal{A}) = \Pr[\mathbf{Exec}^{\mathsf{intsfctxt}}_{\mathsf{AE}}(\mathcal{A}) = 1].$$

We define the notion of existential unforgeablity under chosen message attack of a signature scheme as follows:

Definition 15 (EUF-CMA). *Consider the signature scheme* $\{\mathsf{keysig}, \mathsf{sig}, \mathsf{ver}\}$*, where* keysig *be the key generation method for this scheme. Let* $\mathcal{A}$ *be an adversary that has access to the oracle* $\mathsf{sig}_{sk}(\cdot)$*. We define the experiment as follows:*

$\mathbf{Exec}^{\mathsf{eufcma}}_{(\mathsf{sig},\mathsf{ver})}(\mathcal{A})$
$(pk, sk) \xleftarrow{r} \mathsf{keysig}$
$(m, \sigma) \leftarrow \mathcal{A}^{\mathsf{sig}_{sk}(\cdot)}$
if $\mathsf{ver}_{pk}(m, \sigma) = 1$*; and* m *has not been queried to* $\mathsf{sig}_{sk}(\cdot)$
then return *1* ***else return*** *0*

The attacker's advantage is defined as

$$\mathbf{Adv}^{\mathsf{eufcma}}_{(\mathsf{sig},\mathsf{ver})}(\mathcal{A}) = \Pr[\mathbf{Exec}^{\mathsf{eufcma}}_{(\mathsf{sig},\mathsf{ver})}(\mathcal{A}) = 1].$$

B Jager et al.'s Definition of ACCE

Here we present the revised ACCE definition of Jager et al. [10]. In this definition each oracle $\Pi^s_{i,j}$ maintains an additional internal state variable $b^s_{i,j} \xleftarrow{r} \{0, 1\}$ chosen at random at the start of the game. Further to this an oracle $\Pi^s_{i,j}$ maintains variables $(u^s_{i,j}, v^s_{i,j}, c^s_{i,j}, \theta^s_{i,j})$. The states $u^s_{i,j}$ and $v^s_{i,j}$ are counters (initialised to $(0, 0)$) used to ensure that $\mathcal{A}$ cannot submit a ciphertext previously output by Encrypt oracle to the Decrypt oracle. The variable $c^s_{i,j}$ defines the list of ciphertext output by the encryption oracle, where $c^s_{i,j}[u]$ denotes the u-th entry on the list. Finally, $\theta^s_{i,j}$ stores the pair indices (j, t) necessary to define the partner $\Pi^t_{j,i}$ of $\Pi^s_{i,j}$. The two states st_e and st_d are maintained by encryption and decryption operations of the stateful symmetric encryption scheme (each oracle $\Pi^s_{i,j}$ shall maintain a different set of states). As before we let enc and dec be the encryption and decryption algorithms of our symmetric encryption scheme. The adversary $\mathcal{A}$ will be permitted to make the following queries:

- $\mathsf{Send}^{\mathrm{pre}}(\Pi^s_{i,j}, m)$: This is identical to the Send query in the preliminaries section above, except that it replies with $\perp$ if oracle $\Pi^s_{i,j}$ has state $\delta = \texttt{accept}$ (this shall be handled by the decrypt query).
- $\mathsf{Reveal}(\Pi^s_{i,j})$ and $\mathsf{Corrupt}(i)$ are the standard queries for revealing a session key and corrupting a participant.
- $\mathsf{Encrypt}(\Pi^s_{i,j}, m_0, m_1, h)$: takes as input two equal length messages m_0 and m_1 and a header h. If $\Pi^s_{i,j}$ has $\delta \neq \texttt{accept}$ then $\Pi^s_{i,j}$ returns $\perp$. Otherwise it proceeds with encryption as in Figure 4 dependent on the internal state $b^s_{i,j}$.
- $\mathsf{Decrypt}(\Pi^s_{i,j}, c, h)$: takes as input a ciphertext c and a header h. If $\Pi^s_{i,j}$ has $\delta \neq \texttt{accept}$ then $\Pi^s_{i,j}$ returns $\perp$. Otherwise it proceeds with decryption as in Figure 4.

$\underline{\mathsf{Encrypt}(\Pi^s_{i,j}, m_0, m_1, h)}$	$\underline{\mathsf{Decrypt}(\Pi^s_{i,j}, c, h)}$
$(c^{(0)}, st_e^{(0)}) \leftarrow \mathsf{enc}(k^{\rho}_{\mathsf{enc}}, h, m_0)$	$(j, t) := \theta^s_{i,j}$
$(c^{(1)}, st_e^{(1)}) \leftarrow \mathsf{enc}(k^{\rho}_{\mathsf{enc}}, h, m_1)$	$v^s_{i,j} := v^s_{i,j} + 1$
if $c^{(0)} = \perp$ or $c^{(1)} = \perp$ **then return** $\perp$	**if** $b^s_{i,j} = 0$ **then return** $\perp$
$u^s_{i,j} := u^s_{i,j} + 1$	$(m, st_d) \leftarrow \mathsf{dec}(k^{\rho}_{\mathsf{dec}}, h, c, st_d)$
$(c^s_{i,j}[u^s_{i,j}], st_e) := (c^{(b^s_{i,j})}, st_e^{(b^s_{i,j})})$	**if** $v^s_{i,j} > u^t_{j,i}$ or $c \neq c^t_{j,i}[v^s_{i,j}]$, **then** $\mathsf{phase} := 1$
return $c^s_{i,j}[u^s_{i,j}]$	**if** $\mathsf{phase} = 1$ **then return** m

Fig. 4. Encrypt and Decrypt queries

We define the following game $\mathbf{Exec}^{\mathsf{ACCE}}_{\Pi}(\mathcal{A})$ between an adversary $\mathcal{A}$ and challenger $\mathcal{C}$:

1. The challenger $\mathcal{C}$, generates public/secret key pairs for each user $i \in I$ (by running $\mathcal{G}$) and returns the public keys to $\mathcal{A}$.
2. Adversary $\mathcal{A}$ is allowed to make as many $\mathsf{Send}^{\mathrm{pre}}$, Reveal, $\mathsf{Corrupt}$, $\mathsf{Encrypt}$, $\mathsf{Decrypt}$ queries as it likes.
3. Finally $\mathcal{A}$ outputs a triple (i, j, s, b').

We say the adversary $\mathcal{A}$ wins if it outputs $b' = b^s_{i,j}$. In this case the output of $\mathbf{Exec}^{\mathsf{ACCE}}_{\Pi}(\mathcal{A})$ is set to 1. Otherwise the experiment returns 0. Formally we define the advantage of $\mathcal{A}$ as

$$\mathbf{Adv}^{\mathsf{ACCE}}_{\Pi}(\mathcal{A}) = |\Pr[\mathbf{Exec}^{\mathsf{ACCE}}_{\Pi}(\mathcal{A}) = 1] - 1/2| = |\Pr[b' = b^s_{i,j}] - 1/2|.$$

Definition 16 (ACCE). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a (t, ϵ)-**secure ACCE protocol** if for all adversaries $\mathcal{A}$ running in time t the following conditions hold (where $\epsilon = \epsilon_{EA} + \epsilon_{sAE}$):*

1. *(Entity Authentication/EA): There exists with probability at most ϵ_{EA} an oracle $\Pi^s_{i,j}$ such that:*
   - *$\Pi^s_{i,j}$ accepts when $\mathcal{A}$ issues its τ_0-th query with partner j, and*
   - *P_j is uncorrupted with $\tau_0 < \tau_j$ (i.e. at time of accept), and*
   - *$\mathcal{A}$ did not issue a Reveal-query to oracle $\Pi^t_{j,i}$, such that $\Pi^t_{j,i}$ accepted while having a matching conversation to $\Pi^s_{i,j}$ (if such an oracle exists).*
   - *there is no unique oracle $\Pi^t_{j,i}$ such that $\Pi^s_{i,j}$ has a (wire) matching conversation with $\Pi^t_{j,i}$.*
2. *(Secure Channel/sAE): When $\mathcal{A}$ terminates and outputs (i, j, s, b') such that*
   - *$\Pi^s_{i,j}$ accepts when $\mathcal{A}$ issues its τ_0-th query with intended partner j, and*
   - *P_j is uncorrupted with $\tau_0 < \tau_j$ (i.e. at time of accept), and*
   - *$\mathcal{A}$ did not issue a Reveal-query to $\Pi^s_{i,j}$ nor $\Pi^t_{j,i}$ (such that they had a (wire) matching conversation).*

   the advantage is bounded by $\mathbf{Adv}^{\mathsf{ACCE}}_{\Pi}(\mathcal{A}) = |\Pr[b' = b^s_{i,j}] - 1/2| \leq \epsilon_{sAE}$.

C Previous Models for Secure Channels

C.1 Canetti–Krawczyk

The first attempt to combine the notions of secure key exchange and secure channels was made by Canetti and Krawczyk [7]. Here we shall highlight some of the similarities and differences with our new definitions.

Canetti and Krawczyk define a generic network channels protocol built upon a key exchange scheme and two generic functions *send* and *receive*. Here *send* would take as input some application message and output a message for the channel, *receive* would take as input a channel message and output an application message. The functions *Send* and *receive* may only be called after the key-exchange protocol has been completed, as a result [7] does not take into account protocols where their exists a key-confirmation step where messages are sent over the channel using the send functions. Not only does this create problems in defining protocol execution but it means no scheme of this type can be secure in their model. To facilitate a more modular security analysis Canetti and Krawczyk's approach is to first analyse the key-exchange protocol on its own using a notion for session-key security based on that of Bellare and Rogaway [3]. As a result the model is no longer suitable for analysing protocols with a key-confirmation step which uses the establish session key, as this would allow an adversary to trivially break security. In our model we also define a generic channels protocol but we shall consider protocols which have a key confirmation step utilising the session key.

To analyse the protocol as a whole, Canetti and Krawczyk split security into two parts. To be a secure channel protocol, a protocol must be both a secure encryption protocol and a secure authentication protocol. We will also take this approach, as it provides a more general framework. In some situations we may only require an authenticated channel thus having a separate definition for this can prove very useful.

Finally we discuss how Canetti and Krawczyk choose to handle receive (decryption) queries within their security models. To analyse secure encryption protocols they use an indistinguishability based notion, where the adversary is provided access to a left-or-right 'encryption' oracle. As a result an adversary should not be

able to see the output of a receive call for a message for one participant which was previously output by a send call to his partner. The model therefore restricts by stating that if a plaintext output by *receive* was equal to a previous query to *send*, then this is not returned to the adversary. In our model we make a similar restriction but utilise the state of the encryption and decryption schemes to compare the channel messages output at different times during the protocol run. Canetti and Krawczyk justify their restriction by arguing that comparing the channel messages is overly constrained. Consider the header fields of a network protocol. In particular, the *time-to-live* field is decreased at ever router hop when it travels across a network. Therefore when the message is finally delivered it differs from that originally sent, despite the underlying plaintext message remaining the same. We state that our models can be easily extended to consider protocols with these types of header field by considering equivalence classes of channel messages.

C.2 ACCE Definition of Jager et al.

As mentioned in the introduction Jager et al. [11] combine the notions of authenticated key exchange [3, 4] and LHAE security [15] to give a combined notion of secure channel establishment. In Appendix B we present their (revised [10]) definition in detail and in this section we identify some issues with their approach. Our analysis is not concerned with the length hiding properties used by Jager et al. [11] and Paterson et al. [15] so we omit this aspect and consider only stateful authenticated encryption (sAE).

Reveal Queries We begin with what we argue is the main problem with their definition; namely at what point a Reveal query should be permitted. Reveal queries model unintended leakage of session keys from a participant. Security in the presence of Reveal queries then assures that keys which leak from one session do not affect the security of other sessions. Traditionally, Reveal queries are allowed once a participant has accepted a key. In both the new EMV scheme that we consider (cf. Section 2) and TLS (as considered by Jager et al.), the final part of the key-exchange protocol involves a key confirmation step prior to a key being accepted. Here a message encrypted under the newly established session key is used to perform the final authentication of the sender and confirmation of the key. But if a session key is used prior to being accepted it seems logical that a Reveal should therefore be permitted as soon as keys are derived.

In Jager et al.'s definition they assume that "$\kappa \neq \emptyset$ if and only if $\delta = \texttt{accept}$" while in reality TLS has used κ prior to acceptance in order to encrypt both message m_{11} sent from client to server and message m_{13} sent from server to client. If instead, we allow the adverasary to Reveal as soon as a key is derived then we would be able to perform the following "attack":

- The client outputs the encrypted message m_{11}.
- The adversary reveals the client's key (which is allowed, as the client has derived the key).
- The adversary decrypts m_{11} and then re-encrypts it with new randomness, using the revealed key.
- Finally, the adversary forwards the new ciphertext to the server.
- The server accepts since the decrypted plaintext has not changed.

As a result the client and server will no longer have had a matching conversation. This is a requirement of the ACCE security definition and thus, TLS (and similarly EMV) cannot be proved secure with respect to this definition.

We note that Jager et al. [11] issued a revised version of their paper [10] which alters the definition of ACCE to prevent a similar issue with respect to the message m_{13}. In the first part of the definition they give the following additional restriction:

> *"$\mathcal{A}$ did not issue a Reveal-query to oracle Π_j^t, such that Π_j^t accepted while having a matching conversation to Π_i^s (if such an oracle exists)."*

With this restriction (Jager et al.'s description of) TLS can now be proved secured with relation to ACCE when reveals are only permitted once a key is *accepted*. But we stress that this model still fails to consider

attacks of the form which we describe above, when Reveal queries are permitted as soon as a key has been *derived*. The point is that above, the client has been revealed but has yet to reach an accept state and so does *not* violate the new restriction. The adversary succeeds because the server has accepted without having a matching conversation with the client. In our new definition we shall permit reveal queries as soon as keys are derived, thus capturing all forms of this "reveal" attack. However, this does not mean there is an attack against TLS only that TLS has not been proved secure in this stronger security model.

Channel Messages In practice there are two types of messages sent over the wire during secure channel establishment and use. The first type of message that may be observed will be those used to establish the key. These are then followed by (encrypted) messages sent over the newly established secure channel. An adversary observing such a channel will not necessarily know when messages cease to be part of the key-exchange and become those of the secure channel. Let us consider the situation when an adversary tries to imitate a secure channel message. If a key has yet to be accepted then this message will affect the operation of the key exchange protocol.

The definition of Jager et al. allows the adversary to make three different types of query $\mathsf{Send}^{\mathrm{pre}}$, Encrypt and Decrypt each of which deals with a different type of message. $\mathsf{Send}^{\mathrm{pre}}$ is used only for messages sent as part of the key-exchange. The Encrypt and Decrypt operations will always return an error unless a key has been accepted. But in practice an adversary may not know when an oracle reaches an accept state. Consider the situation where an adversary makes a Decrypt call prior to a key being accepted. The input to both $\mathsf{Send}^{\mathrm{pre}}$ and Decrypt should model messages which have been received on the channel. In Jager et al.'s model an error would immediately be returned by the decryption oracle since no key has been accepted but in reality the message would actually interact with the current state of the key-exchange protocol. It is therefore intuitively more apealing to have a single Send operation which handles both the key-exchange and decryption operations depending on the state of the participant.

Thus, to achieve greater generality and mirror practice more effectively we shall resort to only using a single Send query in our model. When calling Send an adversary will specify a message m and an operation `op`. Basic channel operations may include `SendCh` and `ReceiveCh`. Prior to the completion of the key exchange the operation will be ignored and the message will become part of the key-exchange execution. In addition our definition also allows the channel to have other capabilities (operations) such as sign not previously captured by the aforementioned definition.

D Proof of Theorem 1

The proof of this theorem will be accomplished in the following subsections. Before proceeding with the main proof we first examine a related concept of Key Secrecy for a simpler protocol.

D.1 Key Secrecy

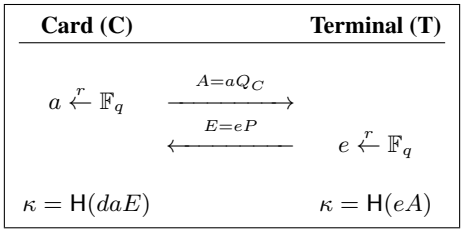

Fig. 5. Unauthenticated Key-Agreement Scheme

We begin our analysis by studying the simpler protocol, π, described in Figure 5. To analyse this protocol we are only interested in whether the secret key remains secret, and so we introduce a new security game to model this fact. Define the following game $\mathbf{Exec}_{\Pi}^{\mathsf{KSec}}(\mathcal{A})$ between an adversary $\mathcal{A}$ and challenger $\mathcal{C}$:

1. The challenger $\mathcal{C}$, generates public/secret key pairs for each user $i \in I$ (by running $\mathcal{G}$) and returns the public keys to $\mathcal{A}$.
2. Adversary $\mathcal{A}$ is allowed to make as many $\mathsf{NewSession}, \mathsf{Send}, \mathsf{Reveal}, \mathsf{Corrupt}$ queries as it likes.
3. Finally $\mathcal{A}$ outputs a pair Π^* and κ^*.

We say the adversary $\mathcal{A}$ wins if Π^* is fresh and κ^* is the key agreed by κ^*. In this case the output of $\mathbf{Exec}_{\Pi}^{\mathsf{Test}}(\mathcal{A})$ is set to 1. Otherwise the output is 0. Formally we define the advantage of $\mathcal{A}$ as

$$\mathbf{Adv}_{\Pi}^{\mathsf{KSec}}(\mathcal{A}) = |\Pr[\mathbf{Exec}_{\Pi}^{\mathsf{KSec}}(\mathcal{A}) = 1]|.$$

Definition 17 (Key Secrecy). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a $(t, \epsilon_{\mathsf{KSec}})$-**key secret AK protocol** if for all adversaries $\mathcal{A}$ running in time t the following holds:*

1. *In the presence of a benign adversary on $\Pi_{i,j}^s$ and $\Pi_{j,i}^t$ both oracles accept holding the same session identifier sid, the same session key κ, and this key is distributed uniformly at random on $\{0,1\}^k$.*
2. *$\mathcal{A}$'s advantage is bounded by $\mathbf{Adv}_{\Pi}^{\mathsf{KSec}}(\mathcal{A}) \leq \epsilon_{\mathsf{KSec}}$.*

We can also define a weaker version of this model for one-sided authentication by running the experiment in the same way as before but changing the winning condition slightly. We say the adversary $\mathcal{A}$ wins the wKSec experiment if Π^* is *OS-fresh* and κ^* is the key agreed by κ^*.

Definition 18 (Weak Key Secrecy). *A protocol $P = \{\Pi, \mathcal{G}\}$ is a $(t, \epsilon_{\mathsf{wKSec}})$-**weak Key-secure AK protocol** if for all adversaries $\mathcal{A}$ running in time t the following holds:*

1. *In the presence of a benign adversary on $\Pi_{i,j}^s$ and $\Pi_{j,i}^t$ both oracles accept holding the same session identifier sid, the same session key κ and this key is distributed uniformly at random on $\{0,1\}^k$.*
2. *$\mathcal{A}$'s advantage is bounded by $\mathbf{Adv}_{\Pi}^{\mathsf{wKSec}}(\mathcal{A}) \leq \epsilon_{\mathsf{wKSec}}$.*

Given this definition we can now analyse the protocol in Figure 5. The proof relies on the following problem being hard.

Definition 19 (Gap Diffie–Hellman). *Let $\mathcal{O}_{\mathsf{DDH}}$ be an oracle that solves the DDH problem in G, i.e. the oracle takes as input $rP, sP, tP \in G$, and outputs one if $tP = rsP$ and zero otherwise.*
The Gap Diffie–Hellman problem then asks that given $aP, bP \in G$ where $a, b \xleftarrow{r} \mathbb{F}_q$, and access to $\mathcal{O}_{\mathsf{DDH}}$, compute abP (i.e. solve CDH). The advantage of an adversary $\mathcal{A}$ against the Gap Diffie–Hellman problem is defined by

$$\mathbf{Adv}_{G}^{\mathsf{Gap\text{-}DH}}(\mathcal{A}) = \Pr[a, b \xleftarrow{r} \mathbb{F}_q : \mathcal{A}^{\mathcal{O}_{\mathsf{DDH}}}(aP, bP) = abP].$$

Lemma 1. *The weak key secrecy of the reduced protocol π is reducible to the Gap Diffie–Hellman assumption, i.e. we have for all adversaries $\mathcal{A}$ there exists an adversary $\mathcal{B}$ such that*

$$\mathbf{Adv}_{\pi}^{\mathsf{wKSec}}(\mathcal{A}) \leq n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}_{E(\mathbb{F}_p)}^{\mathsf{Gap\text{-}DH}}(\mathcal{B}),$$

where n_C is the number of cards in the system and $|h|$ is the output length of the hash function.

Proof. The proof of this lemma uses the technique first presented in [12] for analysing a hashed Diffie–Helman based key agreement protocol. Assume we have an adversary $\mathcal{A}$ against the key secrecy of π we shall use this to construct an adversary $\mathcal{B}$ against Gap Diffie–Hellman, where $\mathcal{B}$ is given the challenge aP, bP.

The algorithm $\mathcal{B}$ begins by setting up n_C authenticated participants by choosing a secret key $d_i \xleftarrow{r} \mathbb{F}_q$ for each authenticated participant $i \in C$ and sets the public key $Q_i = d_i P$ except for one participant $i^* \in C$ where we set the public key to aP. $\mathcal{B}$ also sets up n_T unauthenticated participants.

Algorithm $\mathcal{B}$ will then use its DDH oracle $\mathcal{O}_{DDH}$ to provide simulations of $\mathcal{A}$'s oracles as follows:

– $\mathsf{NewSession}(i, \rho)$ – $\mathcal{B}$ starts a new session for i. All participants may have a total of n_s sessions.
– $\mathsf{Send}(\pi^s_{i,j}, m)$ –
  • For $i \in C$ (and $\rho = \texttt{initiator}$), select at random $\alpha^s_{i,j} \xleftarrow{r} \mathbb{F}_q$ to create message $A = \alpha^s_{i,j} Q_i$.
  • For $i \in T$ ($\rho = \texttt{responder}$), select at random $\beta^s_{i,j} \xleftarrow{r} \mathbb{F}_q$ to create message $E = \beta^s_{i,j} bP$.
  This will result in a shared key $\kappa = \mathsf{H}(\alpha^s_{i^*,j}\beta^s_{i^*,j} abP)$ for oracle $\pi^s_{i^*,j}$ with partner π^t_{j,i^*}, where $j \in T$.
– $\mathsf{Corrupt}(i, d')$ –
  • For $i \in C$, then return d_i and replace it with d' unless $i = i^*$ in which case abort
  • For $i \in T$, return $\perp$.
– $\mathsf{Reveal}(\pi^s_{i,j})$ – To answer Reveal queries, $\mathcal{B}$ will maintain a Guess session key list (G-List). Each element on the G-List is a tuple of the form (τ, i, j, κ_R). Queries are answered as follows:
  • First $\mathcal{B}$ checks the G-list and if there is an entry for i, j then $\mathcal{B}$ outputs the corresponding κ_R.
  • If not then $\mathcal{B}$ checks whether the H-list (see below) contains an (M, h, st_h) with $\mathcal{O}_{DDH}(\alpha^s_{i,j} Q_i, \beta^s_{i,j} bP, M) = 1$. If it does then $\mathcal{B}$ sets $st_h = \{i, j\}$ and adds to G-list (τ, i, j, h).
  • Otherwise $\mathcal{B}$ returns a randomly chosen key.
– $\mathsf{H}(M)$ – To answer hash queries, $\mathcal{B}$ maintains an H-List containing tuples of the form (M, h, st_h). Queries are answered as followed:
  • $\mathcal{B}$ first checks whether M is on the H-list. If it is, then $\mathcal{B}$ outputs h.
  • If not then $\mathcal{B}$ must check whether $\mathsf{H}(M)$ is already an valid entry on the G-list for some pair of participants (i, j) by calling its $\mathcal{O}_{DDH}$.
  • If it is a valid entry for some pair of participants (i, j) then $\mathcal{B}$ returns the corresponding κ_R from the G-list and adds $(M, \kappa_R, \{i, j\})$ to the H-list.
  • Otherwise $\mathcal{B}$ chooses a random hash h and adds (M, h, st_h) to list.

Eventually, $\mathcal{A}$ will output its guess $\pi^* = \pi^s_{i,j}$ and κ^*, The probability that $\mathcal{A}$ chooses $i = i^*$ is $1/n_C$ Note that in this case i^* will not have been corrupted so the simulation has been perfect. At this point $\mathcal{B}$ searches the H-list for the entry $(M^*, \kappa^*, st^*_\kappa)$ corresponding to κ^*, using $\mathcal{O}_{DDH}$ to verify that the entry corresponds to i^*, j. If this entry does not exist then $\mathcal{A}$ must have output a random guess for the key, in which case his probability of success is at best $1/|h|$, where $|h|$ is the size of the output to the function H. Since we assume $\mathcal{A}$ to be a winning adversary with probability $(1 - 1/|h|)$ $\mathcal{A}$ queries H such that his guess is on the H-list. If it is on the list then $\mathcal{B}$ calculates the solution to the gap-DH problem as $(1/\alpha^s_{i^*,j}\beta^s_{i^*,j})M^*$. □

D.2 One-sided Entity Authentication

We now turn to proving the various properties in our main theorem, starting with one-sided entity authentication. We make use of the following definition.

Definition 20 (Computational Diffie–Hellman). *The CDH problem then asks that given $rP, sP \in G$, where $r, s \xleftarrow{r} \mathbb{F}_q$, compute rsP. The advantage of an adversary $\mathcal{A}$ against the CDH problem is defined by*

$$\mathbf{Adv}^{\mathsf{CDH}}_G(\mathcal{A}) = \Pr[r, s \xleftarrow{r} \mathbb{F}_q : \mathcal{A}(rP, sP) = rsP].$$

Lemma 2. *If there exists an adversary $\mathcal{A}$ against Π in the sense of* os-entauth *then there are adversaries $\mathcal{B}, \mathcal{C}, \mathcal{D}$ and $\mathcal{E}$ such that*

$$\begin{aligned}\mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{A}) \leq\ & \mathbf{Adv}^{\mathsf{eufcma}}_{\mathrm{cert}}(\mathcal{B}) + n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}^{\mathsf{Gap\text{-}DH}}_{E(\mathbb{F}_p)}(\mathcal{C}) \\ & + n_S \cdot n_C \cdot \mathbf{Adv}^{\mathsf{intsfctxt-0}}_{\mathsf{AE}}(\mathcal{D}) + n_C^2 \cdot \mathbf{Adv}^{\mathsf{CDH}}_{E(\mathbb{F}_p)}(\mathcal{E}).\end{aligned}$$

where n_C is the number of cards in the system, n_S the number of sessions and $|h|$ is the output size of the hash function.

Proof. Assume we have an adversary $\mathcal{A}$ that wins the os-entauth experiment, that is an oracle $\Pi^t_{j,i}$ with $j \in T$ and $i \in C$ accepts with no os-partner oracle. Let E_{none} be the event that $\Pi^t_{j,i}$ accepts but there exists no oracle $\Pi^s_{i,j}$ with $pid = j$ that has accepted with the same key (and session identifier). Let E_{multi} be the event that $\Pi^t_{j,i}$ accepts but there exists multiple oracles with $pid = j$ that have accepted with the same key (and session identifier). Based on the definition of OS-partnering it is clear that:

$$\Pr[\mathcal{A} \text{ “wins”}] = \Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}}] + \Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{multi}}]$$

Let us analyse these two cases:

(i) $\Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}}]$:
To win in this case an adversary must successfully impersonate a valid card, i.e. there exists $\Pi^s_{i^*,j^*}$ ($i^* \in T$) which accepts with $pid = j^*$ and session identifier sid^*, for which there exists no oracle $\Pi^t_{j^*,i^*}$ with $pid = i^*$ and session identifier sid^*. More specifically $\mathcal{A}$ must send a valid key confirmation message $\mathsf{enc}_{\kappa^*}(cert_{Q^*}||a^*, Q^*; st_e)$. Let F be event that $cert_{Q^*}$ is a valid certificate forgery, i.e. it was not output by the the sig algorithm during the setup phase.

$$\Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}}] = \Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}} \wedge F] + \Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}} \wedge \neg F]$$

Consider each term in turn:

(a) $\Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}} \wedge F]$:
We shall use $\mathcal{A}$ to construct an adversary $\mathcal{B}$ against the EUF-CMA property of the signature scheme that the card issuer used to sign the certificate.
Algorithm $\mathcal{B}$ begins by setting up n_C authenticated participants by choosing secret keys $d_i \in \mathbb{F}_q$ and sets the public key to be $Q_i = d_i P$. Additionally $\mathcal{B}$ calls his sign oracle to generate the certificates for these public keys. $\mathcal{B}$ also sets up n_T unauthenticated participants. $\mathcal{B}$ models $\mathcal{A}$'s NewSession, Reveal, Corrupt, Send queries appropriately using the key material he has generated.
At some point $\mathcal{A}$ issues a query $\mathsf{Send}(\Pi^t_{j^*,i^*}, a^*d^*P, \mathsf{op})$, where d^*P is a “forged” public key for some $j^* \in C$ such that a^* and d^* were chosen by $\mathcal{A}$. $\mathcal{B}$ shall responds by generating the ephemeral public key for terminal i^*, specifically e^*P for some $e^* \in \mathbb{F}_q$. Now $\mathcal{A}$ and the simulated terminal oracle $\Pi^s_{i^*,j^*}$ have derived a key $\kappa^* = \mathsf{H}(a^*d^*e^*P)$. Next in order for $\mathcal{A}$ to get $\Pi^s_{i^*,j^*}$ to accept he must responding with $\mathsf{enc}_{\kappa^*}(\text{cert}_{j^*}, a^*, d^*P)$ such that cert_{j^*} verifies correctly. Upon receipt of $\mathsf{enc}_{\kappa^*}(\text{cert}_{j^*}, a^*, d^*P)$, $\mathcal{B}$ decrypts using κ^* and then outputs $(d^*P, \text{cert}_{j^*})$ as his forgery. Therefore,

$$\Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}} \wedge F] \leq \mathbf{Adv}^{\mathsf{eufcma}}_{\text{cert}}(\mathcal{B}).$$

(b) $\Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}} \wedge \neg F]$:
Let H be the event that $\mathcal{A}$ makes a hash query which reveals the session key.

$$\Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}} \wedge \neg F] = \Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}} \wedge \neg F \wedge H] + \Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}} \wedge \neg F \wedge \neg H]$$

First consider $\Pr[\mathcal{A} \text{ “wins”} \wedge E_{\text{none}} \wedge \neg F \wedge H]$.
We shall use $\mathcal{A}$ to construct an adversary $\mathcal{C}'$ against the wKSec property of the unauthenticated key exchange protocol π (cf. Figure 5). Adversary $\mathcal{C}'$ simulates the environment for $\mathcal{A}$ and begins by calling it's setup algorithm to initialise n_C authenticated participants and n_T unauthenticated participants. $\mathcal{C}'$ models $\mathcal{A}$'s $\mathsf{NewSession}_{\mathcal{A}}$, $\mathsf{Reveal}_{\mathcal{A}}$, $\mathsf{Corrupt}_{\mathcal{A}}$, $\mathsf{Send}_{\mathcal{A}}$ queries appropriately by making the corresponding queries to it's own challenger, i.e. $\mathsf{NewSession}_{\mathcal{C}'}$, $\mathsf{Reveal}_{\mathcal{C}'}$, $\mathsf{Corrupt}_{\mathcal{C}'}$, $\mathsf{Send}_{\mathcal{C}'}$. If $\mathcal{A}$ makes a $\mathsf{Send}_{\mathcal{A}}(\Pi^s_{i,j}, m, \mathsf{op})$ query where $\mathsf{op} = \mathsf{SendCh}$ or $\mathsf{ReceiveCh}$ then $\mathcal{C}'$ will first make a $\mathsf{Reveal}_{\mathcal{C}'}$ query and then performs the necessary encryption or decryption itself. Note that the key revealed will only be forwarded back to $\mathcal{A}$ if he issues the same query to $\mathsf{Reveal}_{\mathcal{A}}$. If $\mathcal{A}$ makes a hash query $\mathsf{H}(m)$, $\mathcal{C}'$ will forward this query to his hash oracle but maintains an H-list of each message and hash, (m, h).

In order to win $\mathcal{A}$ must deduce the session key it has established with $\Pi^s_{i^*,j^*}$ prior to performing any encryption operations, i.e. before the key confirmation step. Therefore, $\mathcal{C}'$ will not have issued a reveal query to $\Pi^s_{i^*,j^*}$ and so the key that $\mathcal{A}$ determines will not violate any of $\mathcal{C}'$'s winning conditions. At some point $\mathcal{A}$ achieves its goal and $\Pi^t_{j^*,i^*}$ accepts. Thus, $\mathcal{A}$ has made a query to H which revealed the session key for $\Pi^t_{j^*,i^*}$. $\mathcal{C}'$ can therefore check which h on his H-list decrypts the confirmation message $\mathcal{A}$ sent to $\Pi^t_{j^*,i^*}$ correctly. $\mathcal{C}'$ outputs $\kappa^* = h$ and $\Pi^t_{j^*,i^*}$. Therefore (by Lemma 1),

$$\Pr[\mathcal{A}\text{ “wins”} \wedge E_{\text{none}} \wedge \neg F \wedge H] \leq \mathbf{Adv}^{\mathsf{wKSec}}_{\pi}(\mathcal{C}') \leq (1 - 1/|h|) \cdot n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}^{\mathsf{Gap\text{-}DH}}_{E(\mathbb{F}_p)}(\mathcal{C}).$$

Next consider $\Pr[\mathcal{A}\text{ “wins”} \wedge E_{\text{none}} \wedge \neg F \wedge \neg H]$.
We shall use $\mathcal{A}$ to construct a new adversary $\mathcal{D}$ against the INT-SFCTXT-0 security of AE, (where INT-SFCTXT-0 is the normal INT-SFCTXT game but the adversary is permitted no encryption queries). $\mathcal{D}$ begins by guessing for which session s^*, card $i^* \in C$ and corresponding terminal $j^* \in T$ he thinks $\mathcal{A}$ will impersonate $i^* \in C$ successfully. What we effectively do is set the output of the random oracle H for the key corresponding to $\Pi^{s^*}_{i^*,j^*}$ to be the key chosen at random for the INT-SFCTXT experiment. All other keys are initialised by $\mathcal{D}$ appropriately. If $\mathcal{A}$ makes a Send query with $\mathtt{op} = \mathtt{SendCh}$ after $\delta = \mathtt{accept}$ for $\Pi^{s^*}_{i^*,j^*}$ then $\mathcal{D}$ shall abort since he is not permitted any encryption queries. All other queries $\mathsf{NewSession}, \mathsf{Reveal}, \mathsf{Corrupt}$ and Send when $\mathtt{op} = \emptyset$ are simulated internally by $\mathcal{D}$ selecting appropriate randomness. Since $\mathcal{A}$ does not make any reveal queries or hash queries corresponding to the key of $\Pi^{s^*}_{i^*,j^*}$ the simulation shall remain perfect. When $\mathcal{A}$ outputs a valid key confirmation message then $\mathcal{D}$ has a valid ciphertext forgery.

$$\Pr[\mathcal{A}\text{ “wins”} \wedge E_{\text{none}} \wedge \neg F \wedge \neg H] \leq n_S n_C \mathbf{Adv}^{\mathsf{intsfctxt}-0}_{\mathsf{AE}}(\mathcal{D}).$$

(ii) $\Pr[\mathcal{A}\text{ “wins”} \wedge E_{\text{multi}}]$:
Here we must consider the case when two card sessions establish the same key with a single terminal.
Let $\mathcal{A}$ be an adversary against the uniqueness of sessions. We shall use $\mathcal{A}$ to construct a new adversary $\mathcal{E}$ that solves the CDH problem given challenge rP, sP.
Algorithm $\mathcal{E}$ begins by setting up n_C authenticated participants by choosing secret keys $d_i \in \mathbb{F}_q$ for each authenticated participant and sets the public keys to be $Q_i = d_i P$. Except for two cards C_1 and C_2 chosen at random (note we also consider the case that $C_1 = C_2$). First $\mathcal{E}$ chooses d, a_1 and a_2 at random from $\mathbb{F}_q$. Next $\mathcal{E}$ sets the public key of C_1 to be $Q_1 = a_1^{-1} drP$ and (when $C_1 \neq C_2$) the public key of C_2 to be $Q_2 = a_2^{-1} dP$.
$\mathcal{E}$ models $\mathcal{A}$'s $\mathsf{NewSession}, \mathsf{Reveal}, \mathsf{Corrupt}, \mathsf{Send}$ queries appropriately using the key material it has generated and necessary randomness. Except for cards C_1 and C_2 where Send queries are modelled such that:

- C_1 first sends $a_1 Q_1 = drP$ to some terminal T_j.
- T_j responds with sP
- The key established between C_1 and T_j is $\mathsf{H}(d(rsP))$. (To model any further send queries with $\mathtt{op} = \mathtt{SendCh}$ or $\mathtt{ReceiveCh}$ $\mathcal{E}$, chooses this hash uniformly at random and uses this as the key to perform the necessary encryptions and decryptions.)
- Next start a new session for C_2 by sending dP. In the case of $C_1 \neq C_2$ this corresponds to $a_2 Q_2$ and in the case of $C_1 = C_2$ this corresponds to $a_2' Q_1$ for some $a_2' \xleftarrow{r} \mathbb{F}_q$.

Finally, adversary $\mathcal{A}$ must impersonate the terminal and send rsP to C_2. This will ensure that C_2 establishes the same session key as the previous session of C_1 and T_j, ($\kappa = \mathsf{H}(drsP)$). The adversary $\mathcal{E}$ then uses $\mathcal{A}$'s impersonated terminal message rsP as its CDH solution.

$$\Pr[\mathcal{A}\text{ “wins”} \wedge E_{\text{multi}}] \leq n_C^2 \mathbf{Adv}^{\mathsf{CDH}}_{E(\mathbb{F}_p)}(\mathcal{E}).$$

□

D.3 One-sided Message Authentication

We now turn to the message authentication property:

Lemma 3. *If the Gap-DH problem is hard in $E(\mathbb{F}_p)$, $\mathsf{AE} = (\mathsf{enc}, \mathsf{dec})$ is an int-sfctxt secure authenticated encryption scheme and Π is secure in the sense of* os-entauth*, then Π is secure in the sense of* os-auth*. In particular if there is an adversary $\mathcal{A}$ against the* os-auth *property then there are adversaries $\mathcal{B}, \mathcal{C}$ and $\mathcal{D}$ such that*

$$\mathbf{Adv}^{\mathsf{os\text{-}auth}}_{\Pi}(\mathcal{A}) \leq n_S \cdot (n_C + n_T) \cdot \mathbf{Adv}^{\mathsf{intsfctxt}}_{\mathsf{AE}}(\mathcal{D}) + n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}^{\mathsf{Gap\text{-}DH}}_{E(\mathbb{F}_p)}(\mathcal{C}) + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}).$$

where n_C is the number of cards in the system, n_T the number of terminals, n_S the number of sessions and $|h|$ is the output size of the hash function.

Proof. We shall prove this result via a sequence of games. Let $\mathcal{A}$ be adversary attacking Π in the sense of auth. **Game 0**: This game is identical to $\mathbf{Exec}^{\mathsf{os\text{-}auth}}_{\Pi}(\mathcal{A})$.

$$\Pr[\mathsf{Game0} \Rightarrow 1] = \mathbf{Adv}^{\mathsf{os\text{-}auth}}_{\Pi}(\mathcal{A})$$

Game 1: This proceeds identically to the previous game but aborts if a terminal ($i \in T$) oracle $\Pi^s_{i,j}$ accepts but has no partner oracle. It is easy to see that

$$\Pr[\mathsf{Game0} \Rightarrow 1] \leq \Pr[\mathsf{Game1} \Rightarrow 1] + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B})$$

Game 2: This proceeds identically to the previous game but aborts if $\mathcal{A}$ makes a query to H which reveals the key for an oracle $\Pi^s_{i,j}$. Again it is easy to see that

$$\Pr[\mathsf{Game1} \Rightarrow 1] \leq \Pr[\mathsf{Game2} \Rightarrow 1] + \mathbf{Adv}^{\mathsf{wKSec}}_{\pi}(\mathcal{C}')$$

Game 3: The challenger now selects at random an oracle $\Pi^{s^*}_{i^*,j^*}$.

The game aborts if $\mathsf{Prefix}(L^{app|rec}_{i,j,s}, L^{app|sen}_{j,i,t}) = \mathsf{false}$ for $(i,j,s) \neq (i^*, j^*, s^*)$. Since i^* is chosen at random from $I = C \cup T$ we have:

$$\Pr[\mathsf{Game2} \Rightarrow 1] \leq n_S \cdot (n_C + n_T) \cdot \Pr[\mathsf{Game3} \Rightarrow 1]$$

It remains to study the probability that $\mathcal{A}$ wins ($\mathsf{Game3} \Rightarrow 1$). We shall use $\mathcal{A}$ in Game3 to construct a new adversary $\mathcal{D}$ against the INT-SFCTXT security of AE. What we effectively do is set the output of the random oracle H for the key corresponding to $\Pi^{s^*}_{i^*,j^*}$ to be the key chosen at random for the INT-SFCTXT experiment. When $\mathcal{A}$ makes a Send query with $\mathtt{op} = \mathtt{SendCh}$ or $\mathtt{ReceiveCh}$ when $\delta = \mathtt{accept}$, $\mathcal{D}$ forwards the message to his enc or dec oracle respectively. All other queries NewSession, Reveal, Corrupt and Send when $\delta \neq \mathtt{accept}$ are simulated internally by $\mathcal{D}$ selecting appropriate randomness. Since $\mathcal{A}$ does not make any reveal queries or hash queries corresponding to the key of $\Pi^{s^*}_{i^*,j^*}$ the simulation shall remain perfect. If $\mathcal{A}$ wins the auth game then $\mathsf{Prefix}(L^{app|rec}_{i^*,j^*,s^*}, L^{app|sen}_{j^*,i^*,t^*}) = \mathsf{false}$ and therefore $\mathcal{A}$ has output a ciphertext forgery which allows $\mathcal{D}$ to win the INT-SFCTXT game. We therefore have,

$$\Pr[\mathsf{Game3} \Rightarrow 1] \leq \mathbf{Adv}^{\mathsf{intsfctxt}}_{\mathsf{AE}}(\mathcal{D})$$

Combining all of the above we obtain

$$\begin{aligned}
\mathbf{Adv}^{\mathsf{os\text{-}auth}}_{\Pi}(\mathcal{A}) &= \Pr[\mathsf{Game0} \Rightarrow 1] \\
&\leq \Pr[\mathsf{Game1} \Rightarrow 1] + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}) \\
&\leq \Pr[\mathsf{Game2} \Rightarrow 1] + \mathbf{Adv}^{\mathsf{wKSec}}_{\pi}(\mathcal{C}') + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}) \\
&\leq n_S \cdot (n_C + n_T) \cdot \Pr[\mathsf{Game3} \Rightarrow 1] + \mathbf{Adv}^{\mathsf{wKSec}}_{\pi}(\mathcal{C}') + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}) \\
&\leq n_S \cdot (n_C + n_T) \cdot \mathbf{Adv}^{\mathsf{intsfctxt}}_{\mathsf{AE}}(\mathcal{D}) + \mathbf{Adv}^{\mathsf{wKSec}}_{\pi}(\mathcal{C}') + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B})
\end{aligned}$$

With the final result following from applying Lemmas 1 and 2. □

D.4 One-sided Message Privacy

We now turn to the message privacy property:

Lemma 4. *If the Gap-DH problem is hard in $E(\mathbb{F}_p)$, $\mathsf{AE} = (\mathsf{enc}, \mathsf{dec})$ is an ind-sfcca secure authenticated-encryption scheme and Π is secure in the sense of* os-entauth*. Then Π is secure in the sense of* os-priv*, i.e. any adversary $\mathcal{A}$ against the* os-priv *property can be turned into adversaries $\mathcal{B}$, $\mathcal{C}$ and $\mathcal{D}$ such that*

$$\mathbf{Adv}^{\mathsf{os\text{-}priv}}_{\Pi}(\mathcal{A}) \leq n_S \cdot (n_C + n_T) \cdot \mathbf{Adv}^{\mathsf{indsfcca}}_{\mathsf{AE}}(\mathcal{D}) + n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}^{\mathsf{Gap\text{-}DH}}_{E(\mathbb{F}_p)}(\mathcal{C}) + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}).$$

Proof. We shall prove this result via a sequence of games. Let $\mathcal{A}$ be adversary attacking Π in the sense of priv.
Game 0: This game is identical to $\mathbf{Exec}^{\mathsf{os\text{-}priv}}_{\Pi}(\mathcal{A})$.

$$\Pr[\mathsf{Game0} \Rightarrow 1] - \frac{1}{2} = \mathbf{Adv}^{\mathsf{os\text{-}priv}}_{\Pi}(\mathcal{A})$$

Game 1: This proceeds identically to the previous game but aborts if a terminal ($i \in T$) oracle $\Pi^s_{i,j}$ accepts but has no partner oracle. It is easy to see that we can define an algorithm $\mathcal{B}'$ such that

$$\Pr[\mathsf{Game0} \Rightarrow 1] \leq \Pr[\mathsf{Game1} \Rightarrow 1] + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B})$$

Game 2: This proceeds identically to the previous game but aborts if $\mathcal{A}$ makes a query to H which reveals the key for an oracle $\Pi^s_{i,j}$. Again it is easy to see that

$$\Pr[\mathsf{Game1} \Rightarrow 1] \leq \Pr[\mathsf{Game2} \Rightarrow 1] + \mathbf{Adv}^{\mathsf{wKSec}}_{\pi}(\mathcal{C}')$$

Game 3: The challenger now selects at random an oracle $\Pi^{s^*}_{i^*,j^*}$. The game aborts if the attacker outputs (i, j, s, b') such that $(i, j, s) \neq (i^*, j^*, s^*)$, the game will instead return a random bit. Since i^* is chosen at random from $I = C \cup T$ we have:

$$\Pr[\mathsf{Game2} \Rightarrow 1] - \frac{1}{2} \leq n_S \cdot (n_C + n_T) \cdot \left(\Pr[\mathsf{Game3} \Rightarrow 1] - \frac{1}{2} \right)$$

It remains to study the probability that $\mathcal{A}$ wins ($\mathsf{Game3} \Rightarrow 1$). We shall use $\mathcal{A}$ in Game3 to construct a new adversary $\mathcal{D}$ against the IND-SFCCA security of AE. What we effectively do is set the output of the random oracle H for the key corresponding to $\Pi^{s^*}_{i^*,j^*}$ to be the key chosen at random for the IND-SFCCA experiment. When $\mathcal{A}$ makes a Send query with $\mathtt{op} = \mathtt{SendCh}$ or $\mathtt{ReceiveCh}$ when $\delta = \mathtt{accept}$, $\mathcal{D}$ forwards the message to his enc or dec oracle respectively. All other queries NewSession, Reveal, Corrupt and Send when $\delta \neq \mathtt{accept}$ are simulated internally by $\mathcal{D}$ selecting appropriate randomness. Since $\mathcal{A}$ does not make any reveal queries or hash queries corresponding to the key of $\Pi^{s^*}_{i^*,j^*}$ the simulation shall remain perfect. When $\mathcal{A}$ outputs its guess (i^*, j^*, s^*, b'), $\mathcal{D}$ shall forward b' as its guess. We therefore have,

$$\Pr[\mathsf{Game3} \Rightarrow 1] - \frac{1}{2} \leq \mathbf{Adv}^{\mathsf{indsfcca}}_{\mathsf{AE}}(\mathcal{D})$$

Combining all of the above, we yield:

$$\begin{aligned}
\mathbf{Adv}^{\mathsf{os\text{-}priv}}_{\Pi}(\mathcal{A}) &= \Pr[\mathsf{Game0} \Rightarrow 1] - \frac{1}{2} \\
&\leq \Pr[\mathsf{Game1} \Rightarrow 1] - \frac{1}{2} + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}) \\
&\leq \Pr[\mathsf{Game2} \Rightarrow 1] - \frac{1}{2} + \mathbf{Adv}^{\mathsf{wKSec}}_{\pi}(\mathcal{C}') + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}) \\
&\leq n_S(n_C + n_T) \left(\Pr[\mathsf{Game3} \Rightarrow 1] - \frac{1}{2} \right) + \mathbf{Adv}^{\mathsf{wKSec}}_{\pi}(\mathcal{C}') + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B}) \\
&\leq n_S(n_C + n_T)\mathbf{Adv}^{\mathsf{indsfcca}}_{\mathsf{AE}}(\mathcal{D}) + \mathbf{Adv}^{\mathsf{wKSec}}_{\pi}(\mathcal{C}') + \mathbf{Adv}^{\mathsf{os\text{-}entauth}}_{\Pi}(\mathcal{B})
\end{aligned}$$

Again the final result follows from applying Lemmas 1 and 2. □

E Proof of Theorem 2

Proof. We shall prove this result via a sequence of games. Let $\mathcal{A}$ be adversary attacking Π in the sense of unlink.

Game 0: This game is identical to $\mathbf{Exec}_{\Pi}^{\mathsf{unlink}}(\mathcal{A})$.

$$\Pr[\mathsf{Game0} \Rightarrow 1] - \frac{1}{2} = \mathbf{Adv}_{\Pi}^{\mathsf{unlink}}(\mathcal{A})$$

Game 1: The challenger now selects at random i_0^* and i_1^*. The game aborts and returns random b' if $\mathcal{A}$ does not output $i_0 = i_0^*$ and $i_1 = i_1^*$. We obtain

$$\Pr[\mathsf{Game0} \Rightarrow 1] - \frac{1}{2} \leq n_C^2 \cdot \left(\Pr[\mathsf{Game1} \Rightarrow 1] - \frac{1}{2}\right)$$

Game 2: This proceeds identically to the previous game but aborts if $\mathcal{A}$ makes a query to H which reveals the key for the oracle $\mathcal{O}$. We obtain

$$\Pr[\mathsf{Game1} \Rightarrow 1] \leq \Pr[\mathsf{Game2} \Rightarrow 1] + \mathbf{Adv}_{\pi}^{\mathsf{wKSec}}(\mathcal{B}')$$

Game 3: This proceeds identically to the previous game except that whenever Send is called with $\mathcal{O}_C$ and $\mathtt{op} = \mathtt{SendCh}$ then the challenger replaces m with a random message which it then encrypts. Again it is easy to see that we obtain

$$\Pr[\mathsf{Game2} \Rightarrow 1] \leq \Pr[\mathsf{Game3} \Rightarrow 1] + \mathbf{Adv}_{\mathsf{AE}}^{\mathsf{indsfcca}}(\mathcal{C})$$

It remains to study the probability that $\mathcal{A}$ wins ($\mathsf{Game2} \Rightarrow 1$). Since ciphertexts are now distributed uniformly at random the only useful information that $\mathcal{A}$ can determine are the public keys $Q_{i_0^*}$, $Q_{i_1^*}$, and the blinded challenge value $aQ_{i_b^*}$. Since a is chosen at random from $\mathbb{F}_q$, then the distributions $(Q_{i_0^*}, Q_{i_1^*}, aQ_{i_0^*})$ and $(Q_{i_0^*}, Q_{i_1^*}, aQ_{i_1^*})$ are identical, i.e. the advantage is zero even if the adversary is computationally unbounded. We therefore have:

$$\Pr[\mathsf{Game3} \Rightarrow 1] - \frac{1}{2} = 0$$

Combining all of the above:

$$\begin{aligned}
\mathbf{Adv}_{\Pi}^{\mathsf{unlink}}(\mathcal{A}) &= \Pr[\mathsf{Game0} \Rightarrow 1] - \frac{1}{2} \\
&\leq n_C^2 \cdot \left(\Pr[\mathsf{Game1} \Rightarrow 1] - \frac{1}{2}\right) \\
&\leq n_C^2 \cdot \left(\Pr[\mathsf{Game2} \Rightarrow 1] - \frac{1}{2} + \mathbf{Adv}_{\pi}^{\mathsf{wKSec}}(\mathcal{B}')\right) \\
&\leq n_C^2 \cdot \left(\Pr[\mathsf{Game3} \Rightarrow 1] - \frac{1}{2} + \mathbf{Adv}_{\mathsf{AE}}^{\mathsf{indsfcca}}(\mathcal{C}) + \mathbf{Adv}_{\pi}^{\mathsf{wKSec}}(\mathcal{B}')\right) \\
&\leq n_C^2 \cdot \left(\mathbf{Adv}_{\mathsf{AE}}^{\mathsf{indsfcca}}(\mathcal{C}) + \mathbf{Adv}_{\pi}^{\mathsf{wKSec}}(\mathcal{B}')\right) \\
&\leq n_C^2 \cdot \left(\mathbf{Adv}_{\mathsf{AE}}^{\mathsf{indsfcca}}(\mathcal{C}) + n_C \cdot (1 - 1/|h|) \cdot \mathbf{Adv}_{E(\mathbb{F}_p)}^{\mathsf{Gap\text{-}DH}}(\mathcal{B})\right)
\end{aligned}$$

□

We note that if we were permitted to have a small (as would be the case in the original EMV proposal) distinguishing the two distributions $(Q_{i_0^*}, Q_{i_1^*}, aQ_{i_0^*})$ and $(Q_{i_0^*}, Q_{i_1^*}, aQ_{i_1^*})$ may no longer be hard. Let l denote the maximum bit length of a. The real question of interest would then be how small can l be before the above problem becomes easy for computationally bounded adversaries. It is clear that the best attack against the problem for $2^l \ll q$ will be Pollard Lambda method [16], which runs in time $O(2^{l/2})$. This implies that a 32-bit randomizer a only gives 16-bits of security and an 80-bit randomizer only gives 40 bits of security.